NEC

N8406-022A 1Gb Intelligent L2 Switch

Application Guide

Part number: 856-126757-106-00
First edition: July 2008

PN# 456-01764-000

Legal notices

© 2008 NEC Corporation.

The information contained herein is subject to change without notice. The only warranties for NEC products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. NEC shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

SunOS™ and Solaris™ are trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Cisco® is a registered trademark of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Part number: 856-126757-106-00

First edition: July 2008

Contents

Accessing the switch
- Introduction ... 6
- Additional references ... 6
- Typographical conventions ... 6
- Management Network ... 7
  - Connecting through the console port ... 7
  - Connecting through Telnet ... 7
  - Connecting through Secure Shell ... 7
- Using the command line interfaces ... 8
  - Configuring an IP interface ... 8
- Using the Browser-based Interface ... 9
- Using Simple Network Management Protocol ... 9
  - SNMP v1.0 ... 9
  - SNMP v3.0 ... 10
  - Default configuration ... 10
  - User configuration ... 10
  - View based configurations ... 11
  - Configuring SNMP trap hosts ... 12
- Secure access to the switch ... 14
  - Setting allowable source IP address ranges ... 14
  - RADIUS authentication and authorization ... 14
  - TACACS+ authentication ... 19
  - Secure Shell and Secure Copy ... 24
- User access control ... 28
  - Setting up user IDs ... 28

Ports and trunking
- Introduction ... 29
- Ports on the switch ... 29
- Port trunk groups ... 29
  - Statistical load distribution ... 30
  - Built-in fault tolerance ... 30
- Before you configure trunks ... 30
- Trunk group configuration rules ... 30
- Port trunking example ... 31
  - Configuring trunk groups (AOS CLI example) ... 32
  - Configuring trunk groups (BBI example) ... 33
- Configurable Trunk Hash algorithm ... 35
- Link Aggregation Control Protocol ... 36
  - Configuring LACP ... 37

VLANs
- Introduction ... 38
- Overview ... 38
- VLANs and port VLAN ID numbers ... 38
  - VLAN numbers ... 38
  - PVID numbers ... 38
  - Viewing and configuring PVIDs ... 39
- VLAN tagging ... 39
- VLANs and IP interfaces ... 42
- VLAN topologies and design considerations ... 42
  - VLAN configuration rules ... 42
- Multiple VLANS with tagging ... 43
  - Configuring the example network ... 44
- FDB static entries ... 51
  - Trunking support for FDB static entries ... 51
  - Configuring a static FDB entry ... 51

Spanning Tree Protocol
- Introduction ... 52
- Overview ... 52
- Bridge Protocol Data Units ... 52

Determining the path for forwarding BPDUs ... 52
Spanning Tree Group configuration guidelines ... 53
Default Spanning Tree configuration ... 53
Adding a VLAN to a Spanning Tree Group ... 53
Creating a VLAN ... 53
Rules for VLAN tagged ports ... 53
Adding and removing ports from STGs ... 54
Assigning cost to ports and trunk groups ... 54
Multiple Spanning Trees ... 54
Why do we need Multiple Spanning Trees? ... 54
VLAN participation in Spanning Tree Groups ... 55
Configuring Multiple Spanning Tree Groups ... 55
Port Fast Forwarding ... 57
Configuring Port Fast Forwarding ... 58
Fast Uplink Convergence ... 58
Configuration guidelines ... 58
Configuring Fast Uplink Convergence ... 58

RSTP and MSTP
Introduction ... 59
Rapid Spanning Tree Protocol ... 59
Port state changes ... 59
Port type and link type ... 59
RSTP configuration guidelines ... 60
RSTP configuration example ... 60
Multiple Spanning Tree Protocol ... 62
MSTP region ... 62
Common Internal Spanning Tree ... 62
MSTP configuration guidelines ... 62
MSTP configuration example ... 62

IGMP Snooping
Introduction ... 67
Overview ... 67
FastLeave ... 67
IGMP Filtering ... 68
Static multicast router ... 68
IGMP Snooping configuration example ... 68

Remote monitoring
Introduction ... 77
Overview ... 77
RMON group 1 — statistics ... 77
RMON group 2 — history ... 80
RMON group 3 — alarms ... 82
RMON group 9 — events ... 86

High availability
Introduction ... 88
Uplink Failure Detection ... 88
Failure Detection Pair ... 89
Spanning Tree Protocol with UFD ... 89
Configuration guidelines ... 89
Monitoring Uplink Failure Detection ... 90
Configuring Uplink Failure Detection ... 90

Troubleshooting tools
Introduction ... 93
Port Mirroring ... 93
Configuring Port Mirroring (AOS CLI example) ... 94
Configuring Port Mirroring (BBI example) ... 95
Other network troubleshooting techniques ... 97
Console and Syslog messages ... 97
Ping ... 97
Trace route ... 97
Statistics and state information ... 97
Customer support tools ... 97

Accessing the switch

Introduction

This guide describes how to use and configure the switch on the Layer2 switch mode. For the information of how to use on the SmartPanel mode, see the SmartPanel Reference Guide. For the information of SSH, RADIUS, and TACACS+ on the SmartPanel mode, this guide will help you.

This guide will help you plan, implement, and administer the switch software. Where possible, each section provides feature overviews, usage examples, and configuration instructions.

- “Accessing the switch” describes how to configure and view information and statistics on the switch over an IP network. This chapter also discusses different methods to manage the switch for remote administrators, such as setting specific IP addresses and using Remote Authentication Dial-in User Service (RADIUS) authentication, Secure Shell (SSH), and Secure Copy (SCP) for secure access to the switch.
- “Ports and port trunking” describes how to group multiple physical ports together to aggregate the bandwidth between large-scale network devices.
- “VLANs” describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs.
- “Spanning Tree Protocol” discusses how spanning trees configure the network so that the switch uses the most efficient path when multiple paths exist.
- “Rapid Spanning Tree Protocol/Multiple Spanning Tree Protocol” describes extensions to the Spanning Tree Protocol that provide rapid convergence of spanning trees for fast reconfiguration of the network.
- “IGMP Snooping” describes how to use IGMP to conserve bandwidth in a multicast-switching environment.
- “Remote Monitoring” describes how to configure the RMON agent on the switch, so the switch can exchange network monitoring data.
- “High Availability” describes how the switch supports high-availability network topologies. This release provides Uplink Failure Detection.
- “Troubleshooting tools” describes Port Mirroring and other troubleshooting techniques.

Additional references

Additional information about installing and configuring the switch is available in the following guides.

- *N8406-022A 1Gb Intelligent L2 Switch User’s Guide*
- *N8406-022A 1Gb Intelligent L2 Switch Command Reference Guide (AOS)*
- *N8406-022A 1Gb Intelligent L2 Switch Command Reference Guide (ISCLI)*
- *N8406-022A 1Gb Intelligent L2 Switch Browser-based Interface Reference Guide*
- *N8406-022A 1Gb Intelligent L2 Switch SmartPanel Reference Guide*

Typographical conventions

The following table describes the typographic styles used in this guide:

Table 1 Typographic conventions

Typeface or symbol	Meaning	Example
`AaBbCc123`	This type depicts onscreen computer output and prompts.	`Main#`
`AaBbCc123`	This type displays in command examples and shows text that must be typed in exactly as shown.	`Main# sys`
`<AaBbCc123>`	This `bracketed` type displays in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets. This also shows guide titles, special terms, or words to be emphasized.	To establish a Telnet session, enter: `host# telnet <IP address>` Read your user guide thoroughly.
`[ ]`	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	`host# ls [-a]`

Management Network

The 1Gb Intelligent L2 Switch is a Switch Module within the Blade Enclosure. The Blade Enclosure includes an Enclosure Manager Card which manages the modules and CPU Blades in the enclosure.

The 1Gb Intelligent L2 Switch communicates with the Enclosure Manager Card through its internal management port (port 19). The factory default settings permit management and control access to the switch through the 10/100 Mbps Ethernet port on the Blade Enclosure, or the built-in console port. You also can use the external Ethernet ports to manage and control this switch.

The switch management network has the following characteristics:

- Port 19 — Management port 19 has the following configuration:
  - Flow control: both
  - Auto-negotiation
  - Untagged
  - Port VLAN ID (PVID): 4095
- VLAN 4095 — Management VLAN 4095 isolates management traffic within the switch. VLAN 4095 contains only one member port (port 19). No other ports can be members of VLAN 4095.
- Interface 256 — Management interface 256 is associated with VLAN 4095. No other interfaces can be associated with VLAN 4095. You can configure the IP address of the management interface manually or through Dynamic Host Control Protocol (DHCP).
- Gateway 4 — This gateway is the default gateway for the management interface.
- STG 32 — If the switch is configured to use multiple spanning trees, spanning tree group 32 (STG 32) contains management VLAN 4095, and no other VLANS are allowed in STG 32. The default status of STG 32 is off.
  If the switch is configured to use Rapid Spanning Tree Protocol, STG 1 contains management VLAN 4095.

To access the switch management interface:

- Use the Enclosure Manager Card internal DHCP server, through Enclosure-Based IP Addressing
- Assign a static IP interface to the switch management interface (interface 256).

Connecting through the console port

Using a null modem cable, you can directly connect to the switch through the console port. A console connection is required in order to configure Telnet or other remote access applications. For more information on establishing console connectivity to the switch, see the *User's Guide*.

Connecting through Telnet

By default, Telnet is enabled on the switch. Once the IP parameters are configured, you can access the CLI from any workstation connected to the network using a Telnet connection. Telnet access provides the same options for a user and an administrator as those available through the console port, minus certain commands. The switch supports four concurrent Telnet connections.

To establish a Telnet connection with the switch, run the Telnet program on your workstation and issue the telnet command, followed by the switch IP address:

```
telnet <switch IP address>
```

Connecting through Secure Shell

By default, the Secure Shell (SSH) protocol is disabled on the switch. SSH enables you to securely log into another computer over a network to execute commands remotely. As a secure alternative to using Telnet to manage switch configuration, SSH ensures that all data sent over the network is encrypted and secure. For more information, see the "Secure Shell and Secure Copy" section later in this chapter. For additional information on the CLI, see the *Command Reference Guide*.

Using the command line interfaces

The command line interface (CLI) can be accessed via local terminal connection or a remote session using Telnet or SSH. The CLI is the most direct method for collecting switch information and performing switch configuration.

The switch provides two CLI modes: The menu-based AOS CLI, and the tree-based ISCLI. You can set the switch to use either CLI mode.

The Main Menu of the AOS CLI, with administrator privileges, is displayed below:

```
[Main Menu]
      info    - Information Menu
      stats   - Statistics Menu
      cfg     - Configuration Menu
      oper    - Operations Command Menu
      boot    - Boot Options Menu
      maint   - Maintenance Menu
      diff    - Show pending config changes  [global command]
      apply   - Apply pending config changes [global command]
      save    - Save updated config to FLASH [global command]
      revert  - Revert pending or applied changes [global command]
      exit    - Exit  [global command, always available]
```

For complete information about the AOS CLI, refer to the *Command Reference Guide (AOS)*.

The ISCLI provides a tree-based command structure, for users familiar with similar products.

An example of a typical ISCLI command is displayed below:

```
Switch(config)# spanning-tree stp 1 enable
```

For complete information about the ISCLI, refer to the *Command Reference Guide (ISCLI)*.

Configuring an IP interface

An IP interface address must be set on the switch to provide management access to the switch over an IP network. By default, the management interface is set up to request its IP address from a DHCP server on the Enclosure Manager Card.

If you configure an IP address manually, the following example shows how to manually configure an IP address on the switch:

1. Configure an IP interface for the Telnet connection, using the sample IP address of 205.21.17.3.
2. The pending subnet mask address and broadcast address are automatically calculated.

```
>> # /cfg/l3/if 256                  (Select IP interface 256)
>> IP Interface 256# addr 205.21.17.3(Assign IP address for the interface)
Current IP address:     0.0.0.0
New pending IP address: 205.21.17.3
Pending new subnet mask: 255.255.255.0
. . . . . .
>> IP Interface 256# ena               (Enable IP interface 256)
```

3. If necessary, configure default gateway.

4. Configuring the default gateways allows the switch to send outbound traffic to the routers.

```
>> IP Interface 256# ../gw 4              (Select default gateway 4)
>> Default gateway 4# addr 205.21.17.1    (Assign IP address for a router)
>> Default gateway 4# ena                 (Enable default gateway 4)
```

5. Apply, verify, and save the configuration.

```
>> Default gateway 4# apply               (Apply the configuration)
>> Default gateway 4# save                (Save the configuration)
>> # /cfg/dump                           (Verify the configuration)
```

NOTE: When the dhcp function on this switch is enabled, the IP address obtained from the DHCP server overrides the static IP address configured manually.

Using the Browser-based Interface

By default, the Browser-based Interface (BBI) protocol is enabled on the switch. The Browser-based Interface (BBI) provides access to the common configuration, management and operation features of the switch through your Web browser. For more information, see the *Browser-based Interface Reference Guide.*

The BBI is organized at a high level as follows:

- Configuration — These menus provide access to the configuration elements for the entire switch.
  - System — Configure general switch configuration elements.
  - Switch ports — Configure switch ports and related features.
  - Port-based port mirroring — Configure mirrored ports and monitoring ports.
  - Layer 2 — Configure Layer 2 features, including trunk groups, VLANs, and Spanning Tree Protocol.
  - RMON menu — Configure Remote Monitoring (RMON) functions.
  - Layer 3 — Configure all of the IP related information, including IGMP Snooping.
  - Uplink Failure Detection — Configure a Failover Pair of Links to Monitor and Links to Disable.
- Statistics — These menus provide access to the switch statistics and state information.
- Dashboard — These menus display settings and operating status of a variety of switch features.

Using Simple Network Management Protocol

The switch software provides SNMP v1.0 and SNMP v3.0 support for access through any network management software.

SNMP v1.0

To access the SNMP agent on the switch, the read and write community strings on the SNMP manager should be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands on the CLI.

```
>> /cfg/sys/ssnmp/rcomm
```

and

```
>> /cfg/sys/ssnmp/wcomm
```

The SNMP manager should be able to reach the management interface or any one of the IP interfaces on the switch.

For the SNMP manager to receive the traps sent out by the SNMP agent on the switch, the trap host on the switch should be configured with the following command:

```
/cfg/sys/ssnmp/snmpv3/taddr
```

For more details, see “Configuring SNMP trap hosts”.

SNMP v3.0

SNMPv3 is an enhanced version of the Simple Network Management Protocol, approved by the Internet Engineering Steering Group in March, 2002. SNMP v3.0 contains additional security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators, and encryption to protect against threats such as masquerade, modification of information, message stream modification, and disclosure.

SNMP v3 ensures that the client can use SNMP v3 to query the MIBs, mainly for security.

To access the SNMP v3.0 menu, enter the following command in the CLI:

```
>> #  /cfg/sys/ssnmp/snmpv3
```

For more information on SNMP MIBs and the commands used to configure SNMP on the switch, see the *Command Reference Guide*.

Default configuration

The switch software has two users by default. Both the users '`adminmd5`' and '`adminsha`' have access to all the MIBs supported by the switch.

1. username 1: `adminmd5/password adminmd5`. Authentication used is MD5.
2. username 2: `adminsha/password adminsha`. Authentication used is SHA.
3. username 3: `v1v2only/password none`.

To configure an SNMP user name, enter the following command from the CLI:

```
>> #  /cfg/sys/ssnmp/snmpv3/usm 6
```

User configuration

Users can be configured to use the authentication/privacy options. Currently we support two authentication algorithms: MD5 and SHA. These can be specified using the command: `/cfg/sys/ssnmp/snmpv3/usm <x>/auth md5|sha`

1. To configure a user with name 'test,' authentication type MD5, and authentication password of 'test,' privacy option DES with privacy password of 'test,' use the following CLI commands:

```
>> # /cfg/sys/ssnmp/snmpv3/usm 5
>> SNMPv3 usmUser 5 # name "test"
>> SNMPv3 usmUser 5 # auth md5
>> SNMPv3 usmUser 5 # authpw test
>> SNMPv3 usmUser 5 # priv des
>> SNMPv3 usmUser 5 # privpw test
```

2. Once a user is configured you need to specify the access level for this user along with the views the user is allowed access to. This is specified in the access table.

```
>> # /cfg/sys/ssnmp/snmpv3/access 5
>> SNMPv3 vacmAccess 5 # name "testgrp"
>> SNMPv3 vacmAccess 5 # level authPriv
>> SNMPv3 vacmAccess 5 # rview "iso"
>> SNMPv3 vacmAccess 5 # wview "iso"
>> SNMPv3 vacmAccess 5 # nview "iso"
```

3. The group table links the user to a particular access group.

```
>> # /cfg/sys/ssnmp/snmpv3/group 5
>> SNMPv3 vacmSecurityToGroup 5 # uname test
>> SNMPv3 vacmSecurityToGroup 5 # gname testgrp
```

If you want to allow user access only to certain MIBs, see the “View based configurations” section.

View based configurations

CLI user equivalent

To configure an SNMP user equivalent to the CLI 'user,' use the following configuration:

```
/c/sys/ssnmp/snmpv3/usm 4
name "usr"
/c/sys/ssnmp/snmpv3/access 3
name "usrgrp"
rview "usr"
wview "usr"
nview "usr"
/c/sys/ssnmp/snmpv3/group 4
uname usr
gname usrgrp
/c/sys/ssnmp/snmpv3/view 6
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.2"
/c/sys/ssnmp/snmpv3/view 7
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.3"
/c/sys/ssnmp/snmpv3/view 8
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.2"
/c/sys/ssnmp/snmpv3/view 9
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.3"
/c/sys/ssnmp/snmpv3/view 10
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.2"
/c/sys/ssnmp/snmpv3/view 11
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.3"
```

CLI oper equivalent

To configure an SNMP user equivalent to the CLI ‘oper’, use the following configuration:

```
/c/sys/ssnmp/snmpv3/usm 5
name "oper"
/c/sys/ssnmp/snmpv3/access 4
name "opergrp"
rview "oper"
wview "oper"
nview "oper"
/c/sys/ssnmp/snmpv3/group 4
uname oper
gname opergrp
/c/sys/ssnmp/snmpv3/view 20
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.2"
/c/sys/ssnmp/snmpv3/view 21
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.3"
/c/sys/ssnmp/snmpv3/view 22
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.2"
/c/sys/ssnmp/snmpv3/view 23
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.3"
/c/sys/ssnmp/snmpv3/view 24
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.2"
/c/sys/ssnmp/snmpv3/view 25
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.3"
```

Configuring SNMP trap hosts

SNMPv1 trap host

1. Configure a user with no authentication and password.

```
/c/sys/ssnmp/snmpv3/usm 10
name "v1trap"
```

2. Configure an access group and group table entries for the user. The command `/c/sys/ssnmp/snmpv3/access <x>/nview` can be used to specify which traps can be received by the user. In the example below the user will receive the traps sent by the switch.

```
/c/sys/ssnmp/snmpv3/access 10
name "v1trap"
model snmpv1
nview "iso"
/c/sys/ssnmp/snmpv3/group 10
model snmpv1
uname v1trap
gname v1trap
```

3. Configure an entry in the notify table.

```
/c/sys/ssnmp/snmpv3/notify 10
name v1trap
tag v1trap
```

4. Specify the IP address and other trap parameters in the targetAddr and targetParam tables. The `c/sys/ssnmp/snmpv3/tparam <x>/uname` command is used to specify the user name used with this targetParam table.

```
/c/sys/ssnmp/snmpv3/taddr 10
name v1trap
addr 47.80.23.245
taglist v1trap
pname v1param
/c/sys/ssnmp/snmpv3/tparam 10
name v1param
mpmodel snmpv1
uname v1trap
model snmpv1
```

5. The community string used in the traps is specified using the community table.

```
/c/sys/ssnmp/snmpv3/comm 10
index v1trap
name public
uname v1trap
```

SNMPv2 trap host configuration

The SNMPv2 trap host configuration is similar to the SNMPv1 trap host configuration. Wherever you specify the model you need to specify `snmpv2` instead of `snmpv1`.

```
c/sys/ssnmp/snmpv3/usm 10
        name "v2trap"
/c/sys/ssnmp/snmpv3/access 10
        name "v2trap"
        model snmpv2
        nview "iso"
/c/sys/ssnmp/snmpv3/group 10
        model snmpv2
        uname v2trap
        gname v2trap
/c/sys/ssnmp/snmpv3/taddr 10
        name v2trap
        addr 47.81.25.66
        taglist v2trap
        pname v2param
/c/sys/ssnmp/snmpv3/tparam 10
        name v2param
        mpmodel snmpv2c
        uname v2trap
        model snmpv2
/c/sys/ssnmp/snmpv3/notify 10
        name v2trap
        tag v2trap
/c/sys/ssnmp/snmpv3/comm 10
        index v2trap
        name public
        uname v2trap
```

SNMPv3 trap host configuration

To configure a user for SNMPv3 traps you can choose to send the traps with both privacy and authentication, with authentication only, or without privacy or authentication.

This is configured in the access table using the command: `/c/sys/ssnmp/snmpv3/access <x>/level`, and `/c/sys/ssnmp/snmpv3/tparam <x>`. The user in the user table should be configured accordingly.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

The following example shows how to configure a SNMPv3 user `v3trap` with authentication only:

```
/c/sys/ssnmp/snmpv3/usm 11
        name "v3trap"
        auth md5
        authpw v3trap
/c/sys/ssnmp/snmpv3/access 11
        name "v3trap"
        level authNoPriv
        nview "iso"
/c/sys/ssnmp/snmpv3/group 11
        uname v3trap
        gname v3trap
/c/sys/ssnmp/snmpv3/taddr 11
        name v3trap
        addr 47.81.25.66
        taglist v3trap
        pname v3param
/c/sys/ssnmp/snmpv3/tparam 11
        name v3param
        uname v3trap
        level authNoPriv
/c/sys/ssnmp/snmpv3/notify 11
        name v3trap
        tag v3trap
```

For more information on using SNMP, see the *Command Reference Guide*.

Secure access to the switch

Secure switch management is needed for environments that perform significant management functions across the Internet. The following are some of the functions for secured management:

- Limiting management users to a specific IP address range. See the “Setting allowable source IP address ranges” section in this chapter.
- Authentication and authorization of remote administrators. See the “RADIUS authentication and authorization” section or the “TACACS+ authentication” section, both later in this chapter.
- Encryption of management information exchanged between the remote administrator and the switch. See the “Secure Shell and Secure Copy” section later in this chapter.

Setting allowable source IP address ranges

To limit access to the switch without having to configure filters for each switch port, you can set a source IP address (or range) that will be allowed to connect to the switch IP interface through Telnet, SSH, SNMP, or the switch browser-based interface (BBI).

When an IP packet reaches the application switch, the source IP address is checked against the range of addresses defined by the management network and management mask. If the source IP address of the host or hosts is within this range, it is allowed to attempt to log in. Any packet addressed to a switch IP interface with a source IP address outside this range is discarded.

Configuring an IP address range for the management network

Configure the management network IP address and mask from the System Menu in the CLI. For example:

```
>> Main# /cfg/sys/access/mgmt/add
Enter Management Network Address: 192.192.192.0
Enter Management Network Mask:  255.255.255.128
```

In this example, the management network is set to 192.192.192.0 and management mask is set to 255.255.255.128. This defines the following range of allowed IP addresses: 192.192.192.1 to 192.192.192.127.

The following source IP addresses are granted or not granted access to the switch:

- A host with a source IP address of 192.192.192.21 falls within the defined range and would be allowed to access the switch.
- A host with a source IP address of 192.192.192.192 falls outside the defined range and is not granted access. To make this source IP address valid, you would need to shift the host to an IP address within the valid range specified by the mnet and mmask or modify the mnet to be 192.192.192.128 and the mmask to be 255.255.255.128. This would put the 192.192.192.192 host within the valid range allowed by the mnet and mmask (192.192.192.128-255).

RADIUS authentication and authorization

The switch supports the Remote Authentication Dial-in User Service (RADIUS) method to authenticate and authorize remote administrators for managing the switch. This method is based on a client/server model. The Remote Access Server (RAS) — the switch — is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes User Datagram Protocol (UDP) over IP, based on Request For Comments (RFC) 2138 and 2866
- A centralized server that stores all the user authorization information
- A client, in this case, the switch

The switch, acting as the RADIUS client, communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFC 2138 and 2866. Transactions between the client and the RADIUS server are authenticated using a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the switch) and the back-end RADIUS server.

How RADIUS authentication works

RADIUS authentication works as follows:

1. A remote administrator connects to the switch and provides the user name and password.
2. Using Authentication/Authorization protocol, the switch sends the request to the authentication server.
3. The authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the switch to grant or deny administrative access.

Configuring RADIUS on the switch (AOS CLI example)

To configure RADIUS on the switch, do the following:

1. Turn RADIUS authentication on, and then configure the Primary and Secondary RADIUS servers. For example:

```
>> Main# /cfg/sys/radius                (Select the RADIUS Server menu)
>> RADIUS Server# on                     (Turn RADIUS on)
Current status: OFF
New status:     ON
>> RADIUS Server# prisrv 10.10.1.1      (Enter primary server IP)
Current primary RADIUS server:     0.0.0.0
New pending primary RADIUS server: 10.10.1.1
>> RADIUS Server# secsrv 10.10.1.2      (Enter secondary server IP)
Current secondary RADIUS server:     0.0.0.0
New pending secondary RADIUS server: 10.10.1.2
```

2. Configure the primary RADIUS secret and secondary RADIUS secret.

```
>> RADIUS Server# secret
Enter new RADIUS secret: <1-32 character secret>
>> RADIUS Server# secret2
Enter new RADIUS second secret: <1-32 character secret>
```

CAUTION: If you configure the RADIUS secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default User Datagram Protocol (UDP) port number used to listen to RADIUS.
4. The well-known port for RADIUS is 1645.

```
>> RADIUS Server# port
Current RADIUS port: 1645
Enter new RADIUS port [1500-3000]: <port number>
```

5. Configure the number of retry attempts for contacting the RADIUS server and the timeout period.

```
>> RADIUS Server# retries
Current RADIUS server retries: 3
Enter new RADIUS server retries [1-3]:<server retries>
>> RADIUS Server# time
Current RADIUS server timeout: 3
Enter new RADIUS server timeout [1-10]: 10  (Enter the timeout period
                                             in seconds)
```

6. Configure the number of retry attempts for contacting the RADIUS server and the timeout period.

```
>> RADIUS Server# apply
>> RADIUS Server# save
```

Configuring RADIUS on the switch (BBI example)

1. Configure RADIUS parameters.
   a. Click the Configure context button.
   b. Open the System folder, and select Radius.

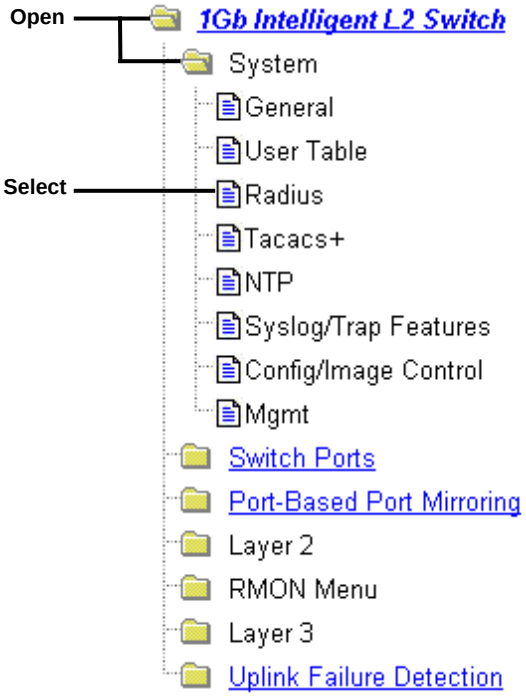

   c. Enter the IP address of the primary and secondary RADIUS servers, and enter the RADIUS secret for each server. Enable the RADIUS server.

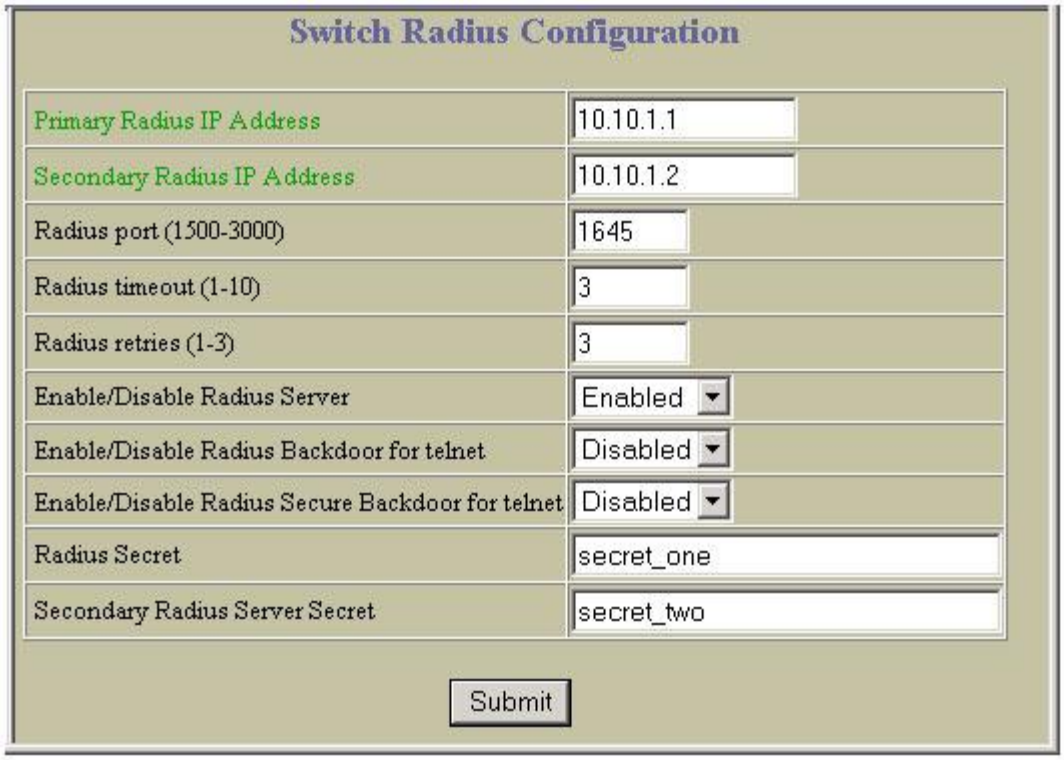

CAUTION: If you configure the RADIUS secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

d. Click Submit.

2. Apply, verify, and save the configuration.

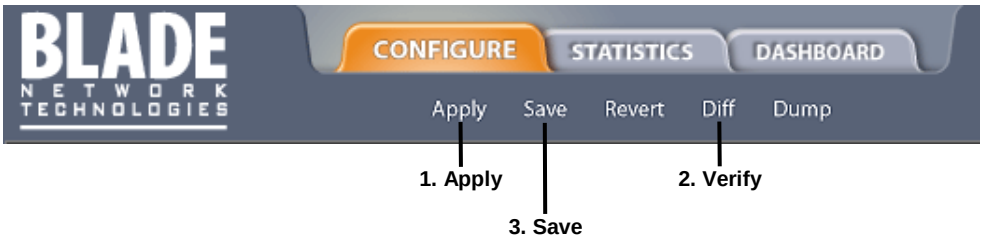

RADIUS authentication features

The switch supports the following RADIUS authentication features:

- Supports RADIUS client on the switch, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows RADIUS secret password up to 32 bytes.
- Supports secondary authentication server so that when the primary authentication server is unreachable, the switch can send client authentication requests to the secondary authentication server. Use the `/cfg/sys/radius/cur` command to show the currently active RADIUS authentication server.
- Supports user-configurable RADIUS server retry and time-out values:
  - Time-out value = 1-10 seconds
  - Retries = 1-3
- The switch will time out if it does not receive a response from the RADIUS server in one to three retries. The switch will also automatically retry connecting to the RADIUS server before it declares the server down.
- Supports user-configurable RADIUS application port. The default is 1645/User Datagram Protocol (UDP)-based on RFC 2138. Port 1812 is also supported.
- Allows network administrator to define privileges for one or more specific users to access the switch at the RADIUS user database.
- Allows the administrator to configure RADIUS backdoor and secure backdoor for Telnet, SSH, HTTP, and HTTPS access.

User accounts for RADIUS users

The user accounts listed in the following table can be defined in the RADIUS server dictionary file.

Table 2 User access levels

User account	Description and tasks performed
User	User interaction with the switch is completely passive; nothing can be changed on the switch. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
Operator	Operators can only effect temporary changes on the switch. These changes are lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation, but do have access to the Maintenance menu. By default, the operator account is disabled and has no password.
Administrator	Administrators are the only ones that can make permanent changes to the switch configuration — changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the switch level. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

RADIUS attributes for user privileges

When the user logs in, the switch authenticates the level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the authentication server successfully authenticates the remote user, the switch verifies the privileges of the remote user and authorizes the appropriate access. The administrator has the option to allow backdoor access through the console port only, or through the console and Telnet/SSH/HTTP/HTTPS access. When backdoor access is enabled, access is allowed even if the primary and secondary authentication servers are reachable. Only when both the primary and secondary authentication servers are not reachable, the administrator has the option to allow secure backdoor (`secbd`) access through the console port only, or through the console and Telnet/SSH/HTTP/HTTPS access. When RADIUS is on, you can have either backdoor or secure backdoor enabled, but not both at the same time. The default value for backdoor access through the console port only is `enabled`. You always can access the switch via the console port, by using `noradius` and the administrator password, whether backdoor/secure backdoor are enabled or not. The default value for backdoor and secure backdoor access through Telnet/SSH/HTTP/HTTPS is `disabled`.

All user privileges, other than those assigned to the administrator, must be defined in the RADIUS dictionary. RADIUS attribute 6, which is built into all RADIUS servers, defines the administrator. The file name of the dictionary is RADIUS vendor-dependent. The RADIUS attributes shown in the following table are defined for user privilege levels.

Table 3 Proprietary attributes for RADIUS

User name/access	User service type	Value
User	Vendor-supplied	255
Operator	Vendor-supplied	252

TACACS+ authentication

The switch software supports authentication, authorization, and accounting with networks using the Cisco Systems TACACS+ protocol. The switch functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the switch either through a data or management port.

TACACS+ offers the following advantages over RADIUS:

- TACACS+ uses TCP-based connection-oriented transport; whereas RADIUS is UDP based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization, and accounting.

How TACACS+ authentication works

TACACS+ works much in the same way as RADIUS authentication.

1. Remote administrator connects to the switch and provides user name and password.

   NOTE: The user name and password can have a maximum length of 128 characters. The password cannot be left blank.

2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. Authentication server checks the request against the user ID database.
4. Using TACACS+ protocol, the authentication server instructs the switch to grant or deny administrative access.

During a session, if additional authorization checking is needed, the switch checks with a TACACS+ server to determine if the user is granted permission to use a particular command.

TACACS+ authentication features

Authentication is the action of determining the identity of a user, and is generally done when the user first attempts to log in to a device or gain access to its services. Switch software supports ASCII inbound login to the device. PAP, CHAP and ARAP login methods, TACACS+ change password requests, and one-time password authentication are not supported.

Authorization

Authorization is the action of determining a user's privileges on the device, and usually takes place after authentication.

The default mapping between TACACS+ authorization privilege levels and switch management access levels is shown in the table below. The privilege levels listed in the following table must be defined on the TACACS+ server.

Table 4 Default TACACS+ privilege levels

User access level	TACACS+ level
user	0
oper	3
admin	6

Alternate mapping between TACACS+ privilege levels and this switch management access levels is shown in the table below. Use the command `/cfg/sys/tacacs/cmap ena` to use the alternate TACACS+ privilege levels.

Table 5 Alternate TACACS+ privilege levels

User access level	TACACS+ level
user	0—1
oper	6— 8
admin	14—15

You can customize the mapping between TACACS+ privilege levels and this switch management access levels. Use the `/cfg/sys/tacacs/usermap` command to manually map each TACACS+ privilege level (0-15) to a corresponding this switch management access level (user, oper, admin, none).

If the remote user is authenticated by the authentication server, the switch verifies the privileges of the remote user and authorizes the appropriate access. When both the primary and secondary authentication servers are not reachable, the administrator has an option to allow backdoor access via the console only or console and Telnet access. The default is disable for Telnet access and enable for console access. The administrator also can enable secure backdoor (/cfg/sys/tacacs/secbd) to allow access if both the primary and secondary TACACS+ servers fail to respond.

Accounting

Accounting is the action of recording a user's activities on the device for the purposes of billing and/or security. It follows the authentication and authorization actions. If the authentication and authorization is not performed via TACACS+, no TACACS+ accounting messages are sent out.

You can use TACACS+ to record and track software logins, configuration changes, and interactive commands.

The switch supports the following TACACS+ accounting attributes:

- protocol (console/telnet/ssh/http)
- start_time
- stop_time
- elapsed_time

NOTE: When using the browser-based Interface, the TACACS+ Accounting Stop records are sent only if the Quit button on the browser is clicked.

Configuring TACACS+ authentication on the switch (AOS CLI example)

1. Turn TACACS+ authentication on, then configure the Primary and Secondary TACACS+ servers.

```
>> Main# /cfg/sys/tacacs              (Select the TACACS+ Server menu)
>> TACACS+ Server# on                 (Turn TACACS+ on)
Current status: OFF
New status: ON
>> TACACS+ Server# prisrv 10.10.1.1   (Enter primary server IP)
Current primary TACACS+ server: 0.0.0.0
New pending primary TACACS+ server: 10.10.1.1
>> TACACS+ Server# secsrv 10.10.1.2   (Enter secondary server IP)
Current secondary TACACS+ server: 0.0.0.0
New pending secondary TACACS+ server: 10.10.1.2
```

2. Configure the TACACS+ secret and second secret.

```
>> TACACS+ Server# secret

Enter new TACACS+ secret: <1-32 character secret>
>> TACACS+ Server# secret2
Enter new TACACS+ second secret: <1-32 character secret>
```

CAUTION: If you configure the TACACS+ secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default TCP port number used to listen to TACACS+.
4. The well-known port for TACACS+ is 49.

```
>> TACACS+ Server# port
Current TACACS+ port: 49
Enter new TACACS+ port [1-65000]: <port number>
```

5. Configure the number retry attempts for contacting the TACACS+ server and the timeout period.

```
>> TACACS+ Server# retries
Current TACACS+ server retries: 3
Enter new TACACS+ server retries [1-3]: 2
>> TACACS+ Server# time
Current TACACS+ server timeout: 5
Enter new TACACS+ server timeout [4-15]: 10 (Enter the timeout period
                                             in minutes)
```

6. Configure custom privilege-level mapping (optional).

```
>> TACACS+ Server# usermap 2
Current privilege mapping for remote privilege 2: not set
Enter new local privilege mapping: user
>> TACACS+ Server# usermap 3 user
>> TACACS+ Server# usermap 4 user
>> TACACS+ Server# usermap 5 oper
```

7. Apply and save the configuration.

Configuring TACACS+ authentication on the switch (BBI example)

1. Configure TACACS+ authentication for the switch.
   a. Click the Configure context button.
   b. Open the System folder, and select Tacacs+.

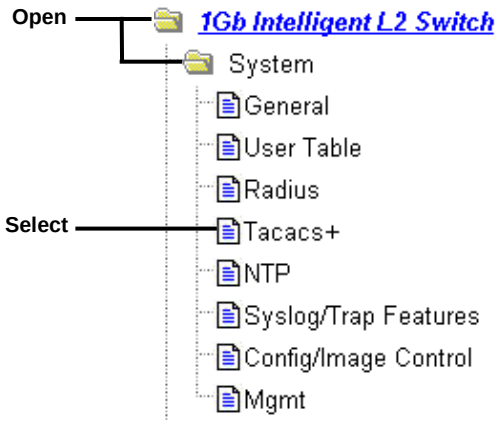

   c. Enter the IP address of the primary and secondary TACACS+ servers, and enter the TACACS+ secret. Enable TACACS+.

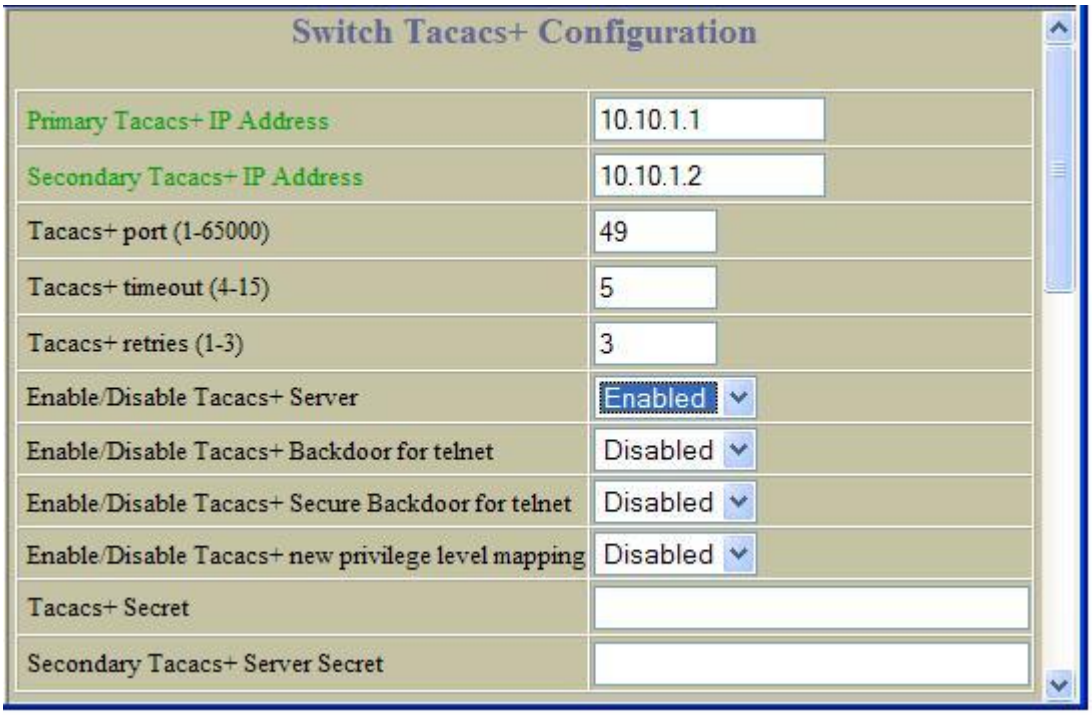

   d. Click Submit.

e. Configure custom privilege-level mapping (optional). Click Submit to accept each mapping change.

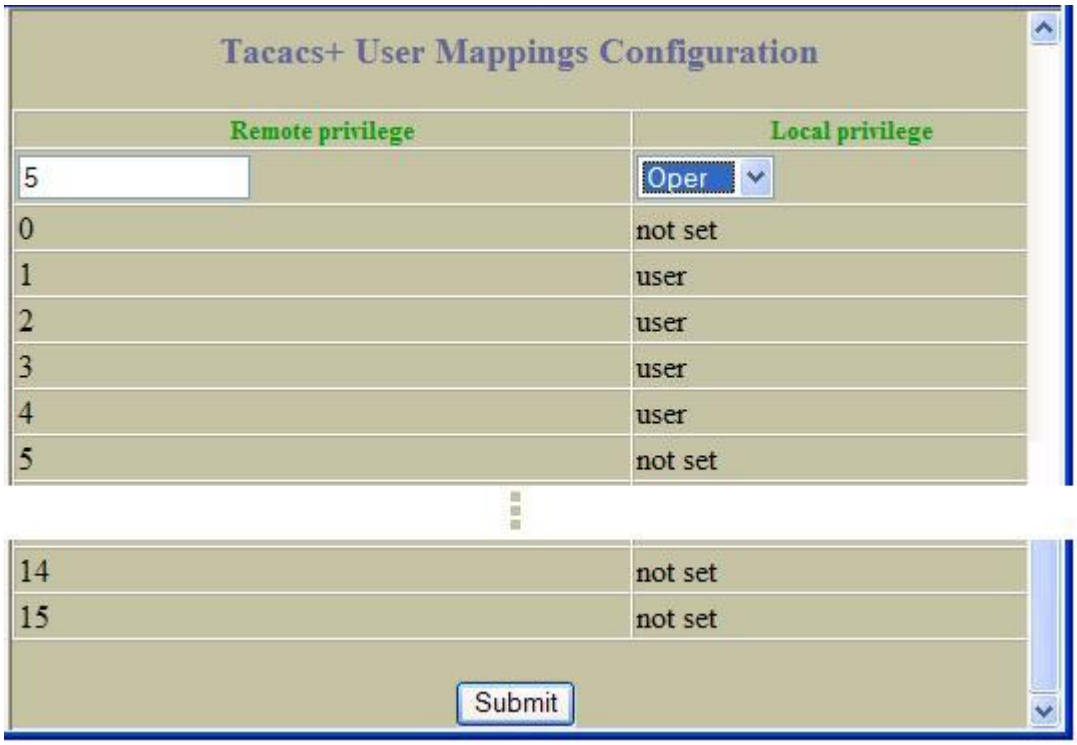

2. Apply, verify, and save the configuration.

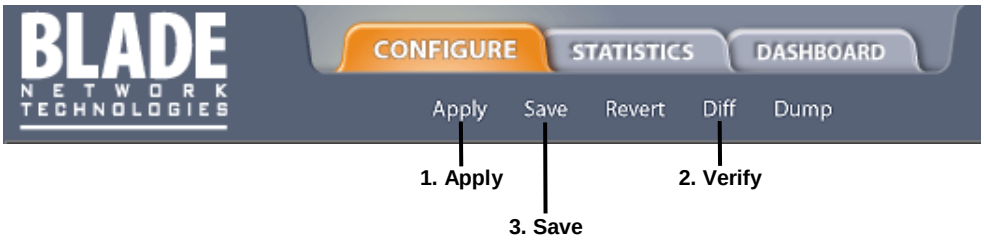

Secure Shell and Secure Copy

Secure Shell (SSH) and Secure Copy (SCP) use secure tunnels to encrypt and secure messages between a remote administrator and the switch. Telnet does not provide this level of security. The Telnet method of managing a switch does not provide a secure connection.

SSH is a protocol that enables remote administrators to log securely into the switch over a network to execute management commands. By default, SSH is disabled (off) on the switch.

SCP is typically used to copy files securely from one machine to another. SCP uses SSH for encryption of data on the network. On a switch, SCP is used to download and upload the switch configuration via secure channels. By default, SCP is disabled on the switch.

The switch implementation of SSH is based on version 1.5 and version 2.0, and supports SSH clients from version 1.0 through version 2.0. Client software can use SSH version 1 or version 2. The following SSH clients are supported:

- SSH 3.0.1 for Linux (freeware)
- SecureCRT® 4.1.8 (VanDyke Technologies, Inc.)
- OpenSSH_3.9 for Linux (FC 3)
- FedoraCore 3 for SCP commands
- PuTTY Release 0.58 (Simon Tatham) for Windows

Configuring SSH and SCP features (AOS CLI example)

Before you can use SSH commands, use the following commands to turn on SSH and SCP.

Enabling or disabling SSH

To enable the SSH feature, connect to the switch CLI and enter the following commands:

```
>> # /cfg/sys/sshd/on            (Turn SSH on)
Current status: OFF
New status: ON
SSHD# apply                       (Apply the changes to start generating
                                   RSA host and server keys)

RSA host key generation starts
. . . . . .
. . . . . .
RSA host key generation completes (lasts 212549 ms)
RSA host key is being saved to Flash ROM, please don’t reboot the box
immediately.
RSA server key generation starts
. . . . . .
RSA server key generation completes (lasts 75503 ms)
RSA server key is being saved to Flash ROM, please don’t reboot the box
immediately.
------------------------------------------------------------------------
Apply complete; don’t forget to “save” updated configuration.
```

NOTE: Secure Shell can be configured using the console port only. SSH menus do not display if you access the switch using Telnet or the Browser-based Interface.

Enabling or disabling SCP apply and save

Enter the following commands from the switch CLI to enable the SCP putcfg_apply and putcfg_apply_save commands:

```
>> # /cfg/sys/sshd/ena        (Enable SCP apply and save)

>> # /cfg/sys/sshd/dis        (Disable SCP apply and save)
SSHD# apply                   (Apply the changes)
```

Configuring the SCP administrator password

To configure the scpadmin (SCP administrator) password, first connect to the switch via the RS-232 management console. For security reasons, the scpadmin password can be configured only when connected directly to the switch console.

To configure the password, enter the following CLI command. At factory default settings, the current SCP administrator password is admin.

```
>> # /cfg/sys/sshd/scpadm
Changing SCP-only Administrator password; validation required. . .
Enter current administrator password: <password>
Enter new SCP-only administrator password: <new password>
Re-enter new SCP-only administrator password: <new password>
New SCP-only administrator password accepted.
```

IMPORTANT: The SCP-only administrator password must be different from the regular administrator password.

Using SSH and SCP client commands

The following shows the format for using some client commands. The examples below use 205.178.15.157 as the IP address of a sample switch.

Logging in to the switch

Enter the following command to log in to the switch:

```
ssh <user>@<switch IP address>
```

For example:

```
>> # ssh admin@205.178.15.157
```

Downloading configuration from the switch using SCP

Enter the following command to download the switch configuration using SCP. You will be prompted for a password:

```
scp <user>@<switch IP address>:getcfg <local filename>
```

For example:

```
>> # scp scpadmin@205.178.15.157:getcfg ad4.cfg
```

The switch prompts you for the scpadmin password.

Uploading configuration to the switch using SCP

Enter the following command to upload configuration to the switch. You will be prompted for a password.

```
scp <local filename> <user>@<switch IP address>:putcfg
```

For example:

```
>> # scp ad4.cfg admin@205.178.15.157:putcfg
```

Applying and saving configuration

Enter the apply and save commands after the command above (scp ad4.cfg 205.178.15.157:putcfg), or use the following commands. You will be prompted for a password.

```
>> # scp <local_filename> <user>@<switch IP addr>:putcfg_apply

>> # scp <local_filename> <user>@<switch IP addr>:putcfg_apply_save
```

For example:

```
>> # scp ad4.cfg admin@205.178.15.157:putcfg_apply

>> # scp ad4.cfg admin@205.178.15.157:putcfg_apply_save
```

Note the following:

- The diff command is automatically executed at the end of putcfg to notify the remote client of the difference between the new and the current configurations.
- putcfg_apply runs the apply command after the putcfg is done.
- putcfg_apply_save saves the new configuration to the flash after putcfg_apply is done.
- The putcfg_apply and putcfg_apply_save commands are provided because extra apply and save commands are usually required after a putcfg.

SSH and SCP encryption of management messages

The following encryption and authentication methods are supported for SSH and SCP:

- Server Host Authentication — Client RSA authenticates the switch at the beginning of every connection
- Key Exchange — RSA
- Encryption — AES256-CBC, AES192-CBC, 3DES-CBC, 3DES, ARCFOUR
- User Authentication — Local password authentication, RADIUS, TACACS+

Generating RSA host and server keys for SSH access

To support the SSH server feature, two sets of RSA keys (host and server keys) are required. The host key is 1024 bits and is used to identify the switch. The server key is 768 bits and is used to make it impossible to decipher a captured session by breaking into the switch at a later time.

When the SSH server is first enabled and applied, the switch automatically generates the RSA host and server keys and is stored in the flash memory.

To configure RSA host and server keys, first connect to the switch console connection (commands are not available via Telnet connection), and enter the following commands to generate them manually:

```
>> # /cfg/sys/sshd/hkeygen                    (Generates the host key)

>> # /cfg/sys/sshd/skeygen                    (Generates the server key)
```

These two commands take effect immediately without the need of an apply command.

When the switch reboots, it will retrieve the host and server keys from the flash memory. If these two keys are not available in the flash memory and if the SSH server feature is enabled, the switch automatically generates them during the system reboot. This process may take several minutes to complete.

The switch can also automatically regenerate the RSA server key. To set the interval of RSA server key autogeneration, use the following command:

```
>> # /cfg/sys/sshd/intrval <number of hours (0-24)>
```

A value of 0 denotes that RSA server key autogeneration is disabled. When greater than 0, the switch will auto generate the RSA server key every specified interval; however, RSA server key generation is skipped if the switch is busy doing other key or cipher generation when the timer expires.

The switch will perform only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the switch is performing key generation at that time, or if another client has logged in immediately prior. Also, key generation will fail if an SSH/SCP client is logging in at that time.

SSH/SCP integration with RADIUS and TACACS+ authentication

SSH/SCP is integrated with RADIUS and TACACS+ authentication. After the RADIUS or TACACS+ server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified RADIUS or TACACS+ servers for authentication. The redirection is transparent to the SSH clients.

User access control

The switch allows an administrator to define end user accounts that permit end users to perform limited actions on the switch. Once end user accounts are configured and enabled, the switch requires username/password authentication.

For example, an administrator can assign a user who can log into the switch and perform operational commands (effective only until the next switch reboot).

The administrator defines access levels for each switch user, as shown in the following table.

Table 6 User access levels

User account	Description	Password
admin	The Administrator has complete access to all menus, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords.	`admin`
oper	The Operator manages all functions of the switch. The Operator can reset ports or the entire switch.	`oper`
user	The User has no direct responsibility for switch management. He/she can view all switch status information and statistics but cannot make any configuration changes to the switch.	`user`

Passwords can be up to 128 characters in length for TACACS+, RADIUS, Telnet, SSH, console, and BBI access.

If RADIUS authentication is used, the user password on the Radius server will override the user password on the switch. Also note that the password-change command on the switch modifies *only* the "use switch" password and has no effect on the user password on the Radius server. RADIUS authentication and user password cannot be used concurrently to access the switch.

Setting up user IDs

The administrator can configure up to 10 user accounts.

To configure an end-user account, perform the following steps:

1. Select a user ID to define.

```
>> # /cfg/sys/access/user/uid 1
```

2. Define the user name and password.

```
>> User ID 1 # name jane            (Assign name "jane" to user ID 1)
Current user name:
New user name: jane
```

3. Define the user access level. By default, the end user is assigned to the user access level. To change the user's access level, enter the user Class of Service (`cos`) command, and select one of the available options.

```
>> User ID 1 # cos <user|oper|admin>
```

4. Enable the user ID.

```
>> # /cfg/sys/access/user/uid <#>/ena
```

Once an end user account is configured and enabled, the user can login to the switch using the username/password combination. The level of switch access is determined by the user CoS for the account. The CoS corresponds to the user access levels described in the User access levels table.

Ports and trunking

Introduction

The first part of this chapter describes the different types of ports used on the switch. This information is useful in understanding other applications described in this guide, from the context of the embedded switch/server environment.

For specific information on how to configure ports for speed, auto-negotiation, and duplex modes, see the port commands in the *Command Reference Guide*.

The second part of this chapter provides configuration background and examples for trunking multiple ports together. Trunk groups can provide super-bandwidth, multi-link connections between switches or other trunk-capable devices. A trunk group is a group of links that act together, combining their bandwidth to create a single, larger virtual link. The switch provides trunking support for the five external ports, two crosslink ports, and 16 server ports.

Ports on the switch

The following table describes the Ethernet ports of the switch, including the port name and function.

NOTE: The actual mapping of switch ports to NIC interfaces is dependant on the operating system software, the type of server blade, and the enclosure type. For more information, see the *User's Guide*.

Table 7 Ethernet switch port names

Port number	Port alias
1	Downlink1
2	Downlink2
3	Downlink3
4	Downlink4
5	Downlink5
6	Downlink6
7	Downlink7
8	Downlink8
9	Downlink9
10	Downlink10
11	Downlink11
12	Downlink12
13	Downlink13
14	Downlink14
15	Downlink15
16	Downlink16
17	XConnect1
18	XConnect2
19	Mgmt
20	Uplink1
21	Uplink2
22	Uplink3
23	Uplink4
24	Uplink5

Port trunk groups

When using port trunk groups between two switches, you can create an aggregate link operating at up to five Gigabits per second, depending on how many physical ports are combined. The switch supports up to 12 trunk groups per switch, each with up to six ports per trunk group.

The trunking software detects broken trunk links (link down or disabled) and redirects traffic to other trunk members within that trunk group. You can only use trunking if each link has the same configuration for speed, flow control, and auto-negotiation.

Statistical load distribution

In a configured trunk group containing more than one port, the load distribution is determined by information embedded within the data frame. For IP traffic, the switch will calculate the trunk port to use for forwarding traffic by implementing the load distribution algorithm on value equals to modulus of (XOR of last 3 bits of Source and last 3 bits of Destination IP address). For non-IP traffic, the switch will calculate the trunk port to use for forwarding traffic by implementing the load distribution algorithm on value equals to modulus of (XOR of last 3 bits of Source and last 3 bits of Destination MAC address).

Built-in fault tolerance

Since each trunk group is composed of multiple physical links, the trunk group is inherently fault tolerant. As long as even one physical link between the switches is available, the trunk remains active.

Statistical load distribution is maintained whenever a link in a trunk group is lost or returned to service.

Before you configure trunks

When you create and enable a trunk, the trunk members (switch ports) take on certain settings necessary for correct operation of the trunking feature.

Before you configure your trunk, you must consider these settings, along with specific configuration rules, as follows:

1. Read the configuration rules provided in the “Trunk group configuration rules” section.
2. Determine which switch ports (up to six) are to become trunk members (the specific ports making up the trunk).
3. Ensure that the chosen switch ports are set to enabled, using the `/cfg/port` command.
4. Trunk member ports must have the same VLAN configuration.
5. Consider how the existing spanning tree will react to the new trunk configuration. See the “Spanning Tree Protocol” chapter for spanning tree group configuration guidelines.
6. Consider how existing VLANs will be affected by the addition of a trunk.

Trunk group configuration rules

The trunking feature operates according to specific configuration rules. When creating trunks, consider the following rules that determine how a trunk group reacts in any network topology:

- All trunks must originate from one device, and lead to one destination device. For example, you cannot combine a link from Server 1 and a link from Server 2 into one trunk group.
- Any physical switch port can belong to only one trunk group.
- Trunking must comply with Cisco® EtherChannel® technology.
- All trunk member ports must be assigned to the same VLAN configuration before the trunk can be enabled.
- All trunk member ports must be set to full duplex mode.
- All trunk member ports must be configured for the same speed.
- If you change the VLAN settings of any trunk member, you cannot apply the change until you change the VLAN settings of all trunk members.
- When an active port is configured in a trunk, the port becomes a trunk member when you enable the trunk using the `/cfg/l2/trunk x/ena` command. The spanning tree parameters for the port then change to reflect the new trunk settings.
- All trunk members must be in the same spanning tree group and can belong to only one spanning tree group. However if all ports are tagged, then all trunk ports can belong to multiple spanning tree groups.
- When a trunk is enabled, the trunk spanning tree participation setting takes precedence over that of any trunk member.
- You cannot configure a trunk member as a monitor port in a Port Mirroring configuration.
- A monitor port cannot monitor trunks; however, trunk members can be monitored.

Port trunking example

In this example, the Gigabit uplink ports on each switch, and the crosslink ports are configured into a total of five trunk groups: two on each switch, and one trunk group at the crosslink between the two switches. All ports operate at Gigabit Ethernet speed.

NOTE: The actual mapping of switch ports to NIC interfaces is dependant on the operating system software, the type of server blade, and the enclosure type. For more information, see the *User's Guide*.

Figure 1 Port trunk group configuration example

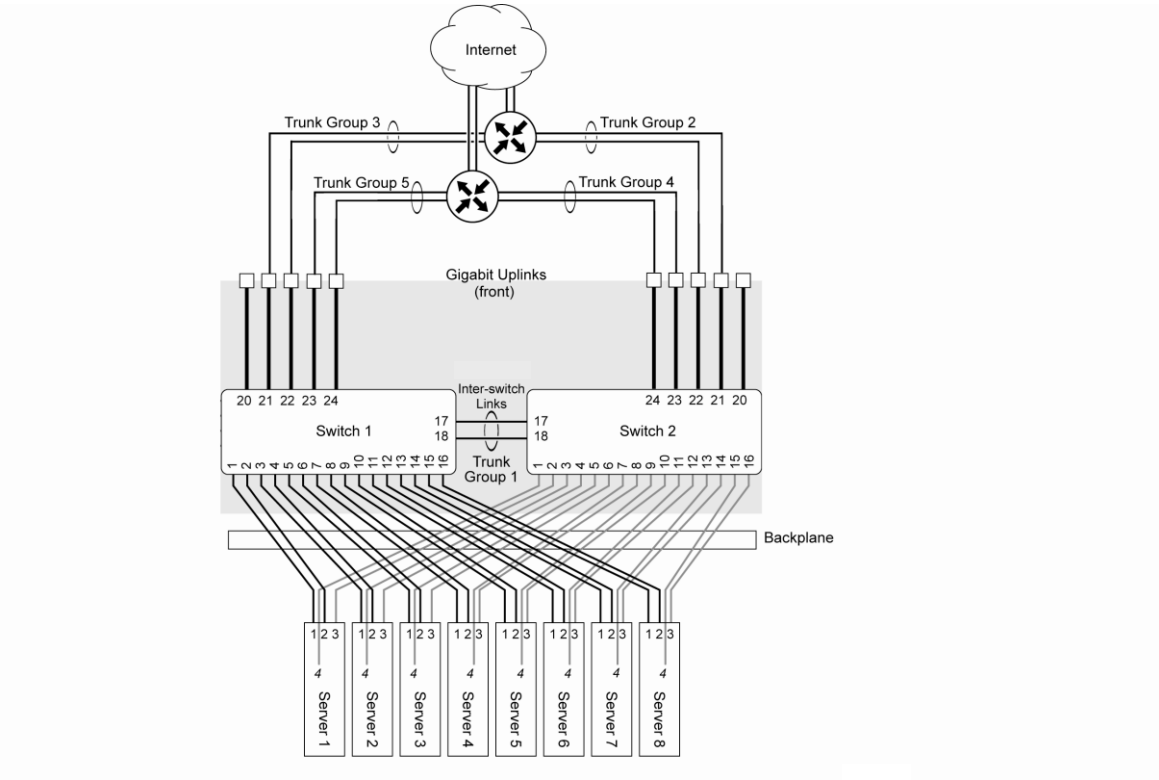

The trunk groups are configured as follows:

- Trunk group 1 is configured by default on the crosslink ports 17 and 18, which connect the switches 1 and 2 together. Since this is the default configuration, you do not need to configure trunk group 1 on either switch. By default, ports 17 and 18 are disabled.
- Trunk groups 2-5 consist of two Gigabit uplink ports each, configured to act as a single link to the upstream routers. The trunk groups on each switch are configured so that there is a link to each router for redundancy.

Prior to configuring each switch in this example, you must connect to the appropriate switch CLI as the administrator. For details about accessing and using any of the commands described in this example, see the *Command Reference Guide*.

Configuring trunk groups (AOS CLI example)

1. On Switch 1, configure trunk groups 5 and 3:

```
>> # /cfg/l2/trunk 5                    (Select trunk group 5)
>> Trunk group 5# add 23                (Add port 23 to trunk group 5)
>> Trunk group 5# add 24                (Add port 24 to trunk group 5)
>> Trunk group 5# ena                   (Enable trunk group 5)
>> Trunk group 5# apply                 (Make your changes active)

>> # /cfg/l2/trunk 3                    (Select trunk group 3)
>> Trunk group 3# add 21                (Add port 21 to trunk group 3)
>> Trunk group 3# add 22                (Add port 22 to trunk group 3)
>> Trunk group 3# ena                   (Enable trunk group 3)
>> Trunk group 3# apply                 (Make your changes active)
>> Trunk group 3# save                  (Save for restore after reboot)
```

2. On Switch 2, configure trunk groups 4 and 2:

```
>> # /cfg/l2/trunk 4                    (Select trunk group 4)
>> Trunk group 4# add 23                (Add port 23 to trunk group 4)
>> Trunk group 4# add 24                (Add port 24 to trunk group 4)
>> Trunk group 4# ena                   (Enable trunk group 4)
>> Trunk group 4# apply                 (Make your changes active)

>> # /cfg/l2/trunk 2                    (Select trunk group 2)
>> Trunk group 2# add 21                (Add port 21 to trunk group 2)
>> Trunk group 2# add 22                (Add port 22 to trunk group 2)
>> Trunk group 2# ena                   (Enable trunk group 2)
>> Trunk group 2# apply                 (Make your changes active)
>> Trunk group 2# save                  (Save for restore after reboot)
```

NOTE: In this example, two switches are used. Any third-party device supporting link aggregation should be configured manually. Connection problems could arise when using automatic trunk group negotiation on the third-party device.

3. Examine the trunking information on each switch using the following command:

```
>> /info/l2/trunk                     (View trunking information)
```

Information about each port in each configured trunk group will be displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state.

Configuring trunk groups (BBI example)

1. Configure trunk groups.
   a. Click the Configure context button on the Toolbar.
   b. Open the Layer 2 folder, and select Trunk Groups.

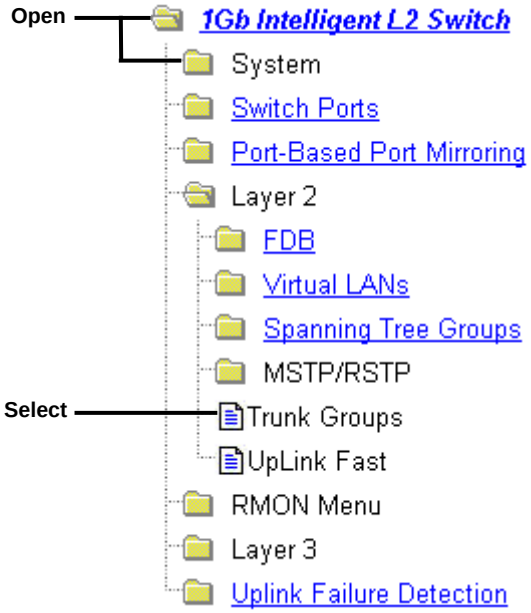

   c. Click a Trunk Group number to select it.

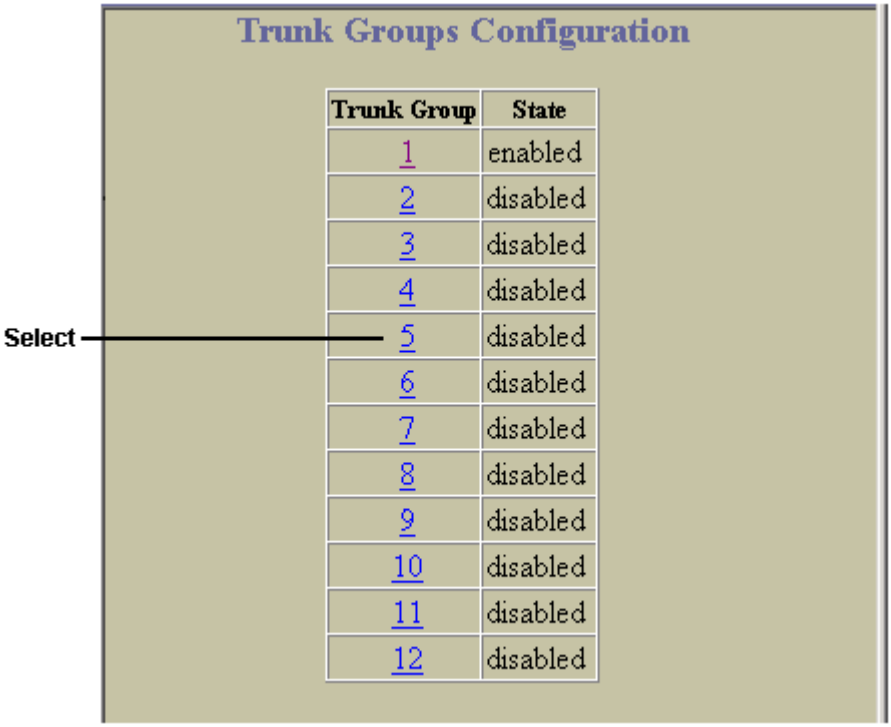

d. Enable the Trunk Group. To add ports, select each port in the Ports Available list, and click Add.

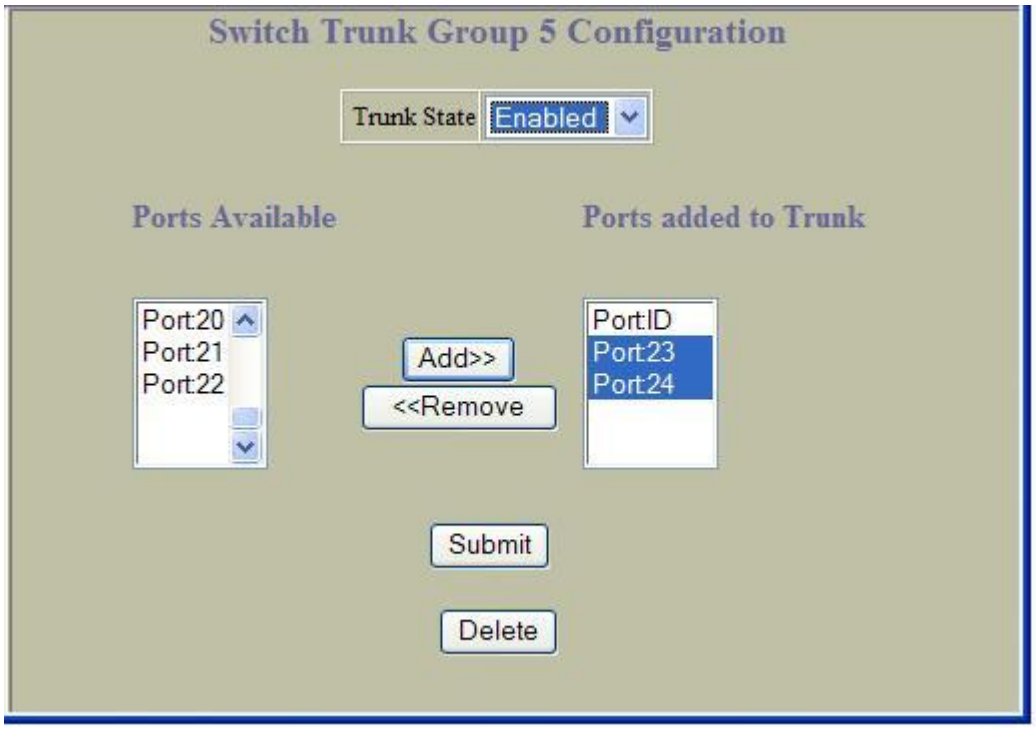

e. Click Submit.

2. Apply, verify, and save the configuration.

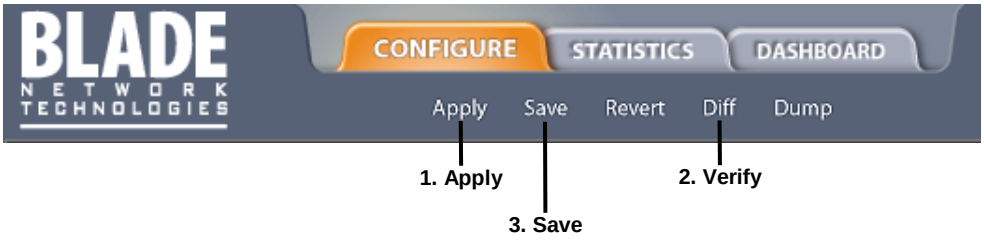

3. Examine the trunking information on each switch.

a. Click the Dashboard context button on the Toolbar.

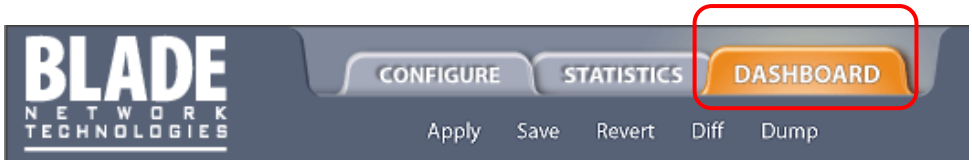

b. Select Trunk Groups.

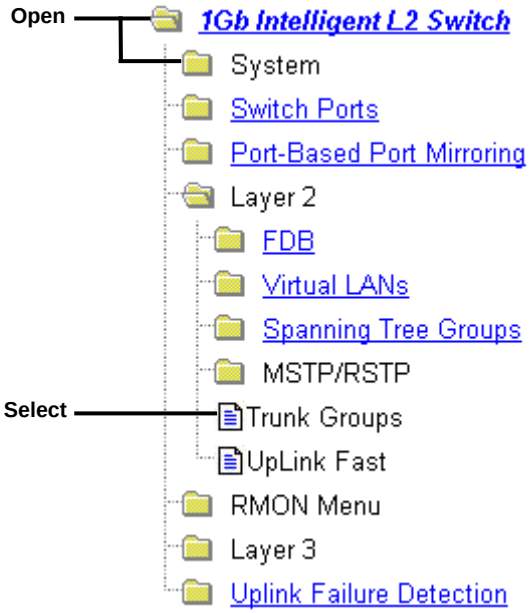

c. Information about each configured trunk group is displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state.

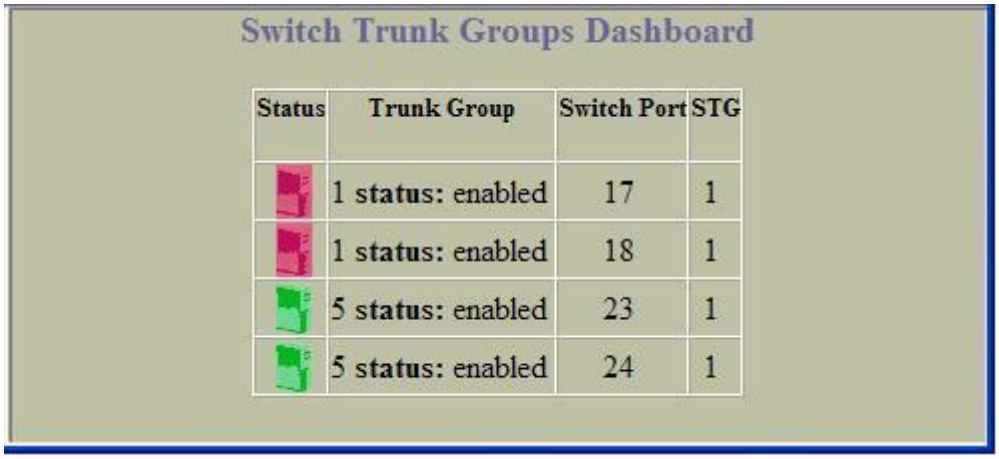

Configurable Trunk Hash algorithm

This feature allows you to configure the particular parameters for the switch Trunk Hash algorithm instead of having to utilize the defaults. You can configure new default behavior for Layer 2 traffic and Layer 3 traffic, using the CLI menu `cfg/l2/thash`. You can select a minimum of one or a maximum of two parameters to create one of the following configurations:

- Source IP (SIP)
- Destination IP (DIP)
- Source MAC (SMAC)
- Destination MAC (DMAC)
- Source IP (SIP) + Destination IP (DIP)
- Source MAC (SMAC) + Destination MAC (DMAC)

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group or Link Aggregation group) with any device that supports the standard. Refer to the IEEE 802.3ad-2002 for a full description of the standard.

The 802.3ad standard allows standard Ethernet links to form a single Layer 2 link using the Link Aggregation Control Protocol (LACP). Link aggregation is a method of grouping physical link segments of the same media type and speed in full duplex, and treating them as if they were part of a single, logical link segment. If a link in a LACP trunk group fails, traffic is reassigned dynamically to the remaining link(s) of the dynamic trunk group.

NOTE: Currently, LACP implementation does not support the Churn machine, an option used to detect if the port is operable within a bounded time period between the actor and the partner. Only the Marker Responder is implemented, and there is no marker protocol generator.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the *system ID* and the port's *admin key*, as follows:

IMPORTANT: System ID—The system ID is an integer value based on the switch's MAC address and the system priority assigned in the CLI.

- Admin key—A port's Admin key is an integer value (1-65535) that you can configure in the CLI. Each switch port that participates in the same LACP trunk group must have the same *admin key* value. The Admin key is *local significant*, which means the partner switch does not need to use the same Admin key value.

For example, consider two switches, an Actor (this switch) and a Partner (another switch), as shown in the following table:

Table 8 Actor vs. partner LACP configuration

Actor Switch	Partner Switch 1	Partner Switch 2
Port 20 (admin key = 100)	Port 1 (admin key = 50)	
Port 21 (admin key = 100)	Port 2 (admin key = 50)	
Port 22 (admin key = 200)		Port 3 (admin key = 60)
Port 23 (admin key = 200)		Port 4 (admin key = 60)

In the configuration shown in the table above, Actor switch ports 20 and 21 aggregate to form an LACP trunk group with Partner switch ports 1 and 2. At the same time, Actor switch ports 22 and 23 form a different LACP trunk group with a different partner.

LACP automatically determines which member links can be aggregated and then aggregates them. It provides for the controlled addition and removal of physical links for the link aggregation.

Each port in the switch can have one of the following LACP modes.

- off (default)—The user can configure this port in to a regular static trunk group.
- active—The port is capable of forming an LACP trunk. This port sends LACPDU packets to partner system ports.
- passive—The port is capable of forming an LACP trunk. This port only responds to the LACPDU packets sent from an LACP *active* port.

Each active LACP port transmits LACP data units (LACPDUs), while each passive LACP port listens for LACPDUs. During LACP negotiation, the admin key is exchanged. The LACP trunk group is enabled as long as the information matches at both ends of the link. If the admin key value changes for a port at either end of the link, that port's association with the LACP trunk group is lost.

When the system is initialized, all ports by default are in LACP *off* mode and are assigned unique *admin keys*. To make a group of ports aggregatable, you assign them all the same *admin key*. You must set the port's LACP mode to *active* to activate LACP negotiation. You can set other port's LACP mode to passive, to reduce the amount of LACPDU traffic at the initial trunk-forming stage.

Use the `/info/l2/trunk` command or the `/info/l2/lacp/dump` command to check whether the ports are trunked.

Configuring LACP

Use the following procedure to configure LACP for port 20 and port 21 to participate in link aggregation.

4. Set the LACP mode on port 20.

```
>> # /cfg/l2/lacp/port 20           (Select port 20)
>> LACP port 20# mode active        (Set port 20 to LACP active mode)
```

5. Define the admin key on port 20. Only ports with the same admin key can form a LACP trunk group.

```
>> LACP port 20# adminkey 100          (Set port 20 adminkey to 100)
Current LACP port adminkey:     20
New pending LACP port adminkey: 100
```

6. Set the LACP mode on port 21.

```
>> # /cfg/l2/lacp/port 21             (Select port 21)
>> LACP port 21# mode active           (Set port 21 to LACP active mode)
```

7. Define the admin key on port 21.

```
>> LACP port 21# adminkey 100          (Set port 21 adminkey to 100)
Current LACP port adminkey:     21
New pending LACP port adminkey: 100
```

8. Apply and verify the configuration.

```
>> LACP port 21# apply              (Make your changes active)
>> LACP port 21# cur                (View current trunking configuration)
```

9. Save your new configuration changes.

```
>> LACP port 21# save               (Save for restore after reboot)
```

VLANs

Introduction

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs). VLANs are commonly used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments.

The following topics are discussed in this chapter:

- VLANs and Port VLAN ID Numbers
- VLAN Tagging
- VLANs and IP Interfaces
- VLAN Topologies and Design Considerations

NOTE: Basic VLANs can be configured during initial switch configuration.

More comprehensive VLAN configuration can be done from the command line interface. See the *Command Reference Guide.*

Overview

Setting up VLANs is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belongs to one broadcast domain.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN.

VLANs and port VLAN ID numbers

VLAN numbers

This switch supports up to 1,000 VLANs per switch. Even though the maximum number of VLANs supported at any given time is 1,000, each can be identified with any number between 1 and 4095. VLAN 1 is the default VLAN, and all ports (except port 19) are assigned to it. VLAN 4095 is reserved for switch management, and it cannot be configured.

Viewing VLANs

The VLAN information menu (/info/l2/vlan) displays all configured VLANs and all member ports that have an active link state, for example:

```
>> Layer 2# vlan
VLAN                Name                 Status Ports
----  -------------------------------- ------ ----------------------
1     Default VLAN                       ena   1 4-18 20-24
2     VLAN 2                             ena   2 3
4095  VLAN 4095                          ena   19
```

PVID numbers

Each port in the switch has a configurable default VLAN number, known as its PVID. This places all ports on the same VLAN initially, although each port PVID is configurable to any VLAN number between 1 and 4094.

The default configuration settings for switches have all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in the following figure, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID =1).

Viewing and configuring PVIDs

You can view PVIDs from the following AOS CLI commands:

Port information

```
>> /info/port
Port Tag RMON PVID      NAME                    VLAN(s)
---- --- ---- ---- -------------- --------------------------------
  1   n   d      1 Downlink1      1
  2   n   e      1 Downlink2      1
  3   n   d      1 Downlink3      1
  4   n   d      1 Downlink4      1
  5   n   d      1 Downlink5      1
  6   n   d      1 Downlink6      1
  7   n   d      1 Downlink7      1
  :
  :
```

Port configuration

```
>>  /cfg/port 22/pvid 22
Current port VLAN ID:     1
New pending port VLAN ID: 22

>> Port 22#
```

Each port on the switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN tagging enabled. See the "VLAN tagging" section in this chapter.

Any untagged frames (those with no VLAN specified) are classified with the PVID of the sending port.

VLAN tagging

The switch supports IEEE 802.1Q VLAN tagging, providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header, allowing each port to belong to multiple VLANs. When you configure multiple VLANs on a port, you must also enable tagging on that port.

Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging is not enabled.

Important terms used with the 802.1Q tagging feature are:

- VLAN identifier (VID) — the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID) — a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID =3) assigns all untagged frames received on this port to VLAN 3.
- Tagged frame — a frame that carries VLAN tagging information in the header. The VLAN tagging information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the switch through a port that is configured as a tagged port.
- Untagged frame — a frame that does not carry any VLAN tagging information in the frame header.
- Untagged member — a port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member — a port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).

NOTE: If an 802.1Q tagged frame is sent to a port that has VLAN-tagging disabled, then the frames are forwarded based on their port-VLAN ID (PVID).

Figure 2 Default VLAN settings

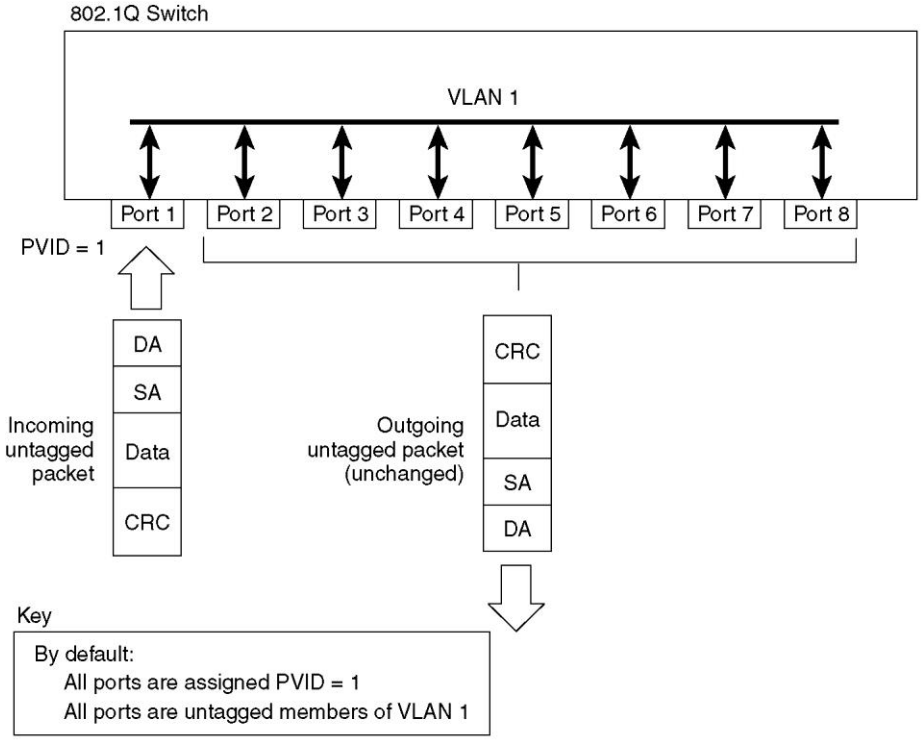

NOTE: The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your switch model.

When you configure VLANs, you configure the switch ports as tagged or untagged members of specific VLANs. See the following figures.

In the following figure, the untagged incoming packet is assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

Figure 3 Port-based VLAN assignment

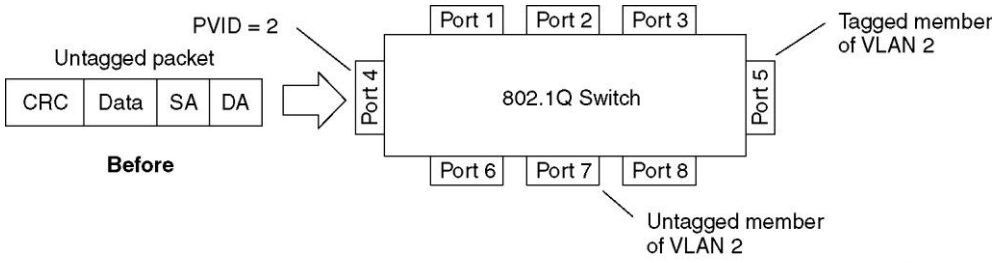

As shown in the following figure, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 4 802.1Q tagging (after port-based VLAN assignment)

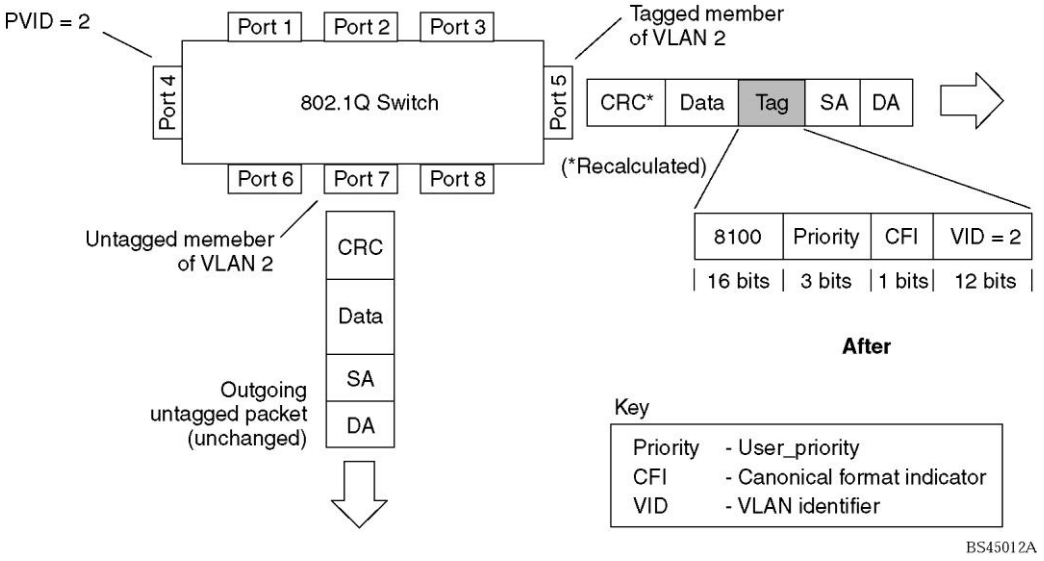

In the following figure, the tagged incoming packet is assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

Figure 5 802.1Q tag assignment

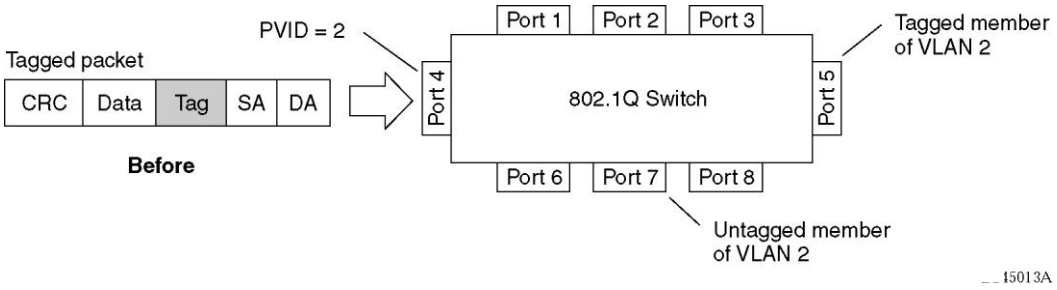

As shown in the following figure, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 6 802.1Q tagging (after 802.1Q tag assignment)

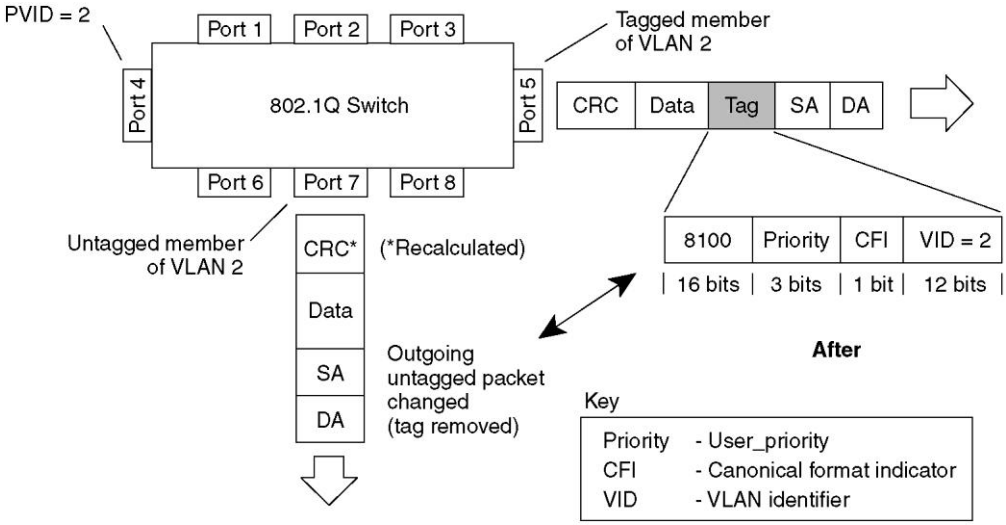

NOTE: Using the `/boot/conf factory` command resets all ports to VLAN 1 (except port 19) and all other settings to the factory defaults at the next reboot.

VLANs and IP interfaces

Carefully consider how you create VLANs within the switch, so that communication with the switch remains possible. In order to access the switch for remote configuration, trap messages, and other management functions, be sure that at least one IP interface on the switch has a VLAN defined.

You can also inadvertently cut off access to management functions if you exclude the ports from the VLAN membership. For example, if all IP interfaces are left on VLAN 1 (the default), and all ports are configured for VLAN 2, and then switch management features are effectively cut off.

To remedy this, keep all ports used for remote switch management on the default VLAN and assign an IP interface to the default VLAN.

For more information on configuring IP interfaces, see the “Configuring an IP interface” section in the “Accessing the switch” chapter.

VLAN topologies and design considerations

By default, all switch ports are configured to the default VLAN 1. This configuration groups all ports into the same broadcast domain. The VLAN has an 802.1Q VLAN ID of 1. VLAN tagging is turned off, because, by default, all ports are members of a single VLAN only.

If configuring Spanning Tree Protocol (`/cfg/l2/stp`), note that each of spanning tree groups 2-32 may contain only one VLAN. If configuring Multiple Spanning Tree Protocol (`/cfg/l2/mrst`), each of spanning tree groups 1-32 may contain multiple VLANs.

VLAN configuration rules

VLANs operate according to specific configuration rules which must be considered when creating VLANs. For example:

- We recommend that all ports involved in trunking and Port Mirroring have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed. For more information on port trunking, see the “Port trunking example” section in the “Ports and trunking” chapter.
- All ports that are involved in Port Mirroring must have memberships in the same VLANs. If a port is configured for Port Mirroring, the port’s VLAN membership cannot be changed. For more information on configuring Port Mirroring, see the “Port Mirroring” section in the “Troubleshooting tools” appendix.
- When you delete a VLAN, untagged ports are moved to the default VLAN (VLAN 1). Tagged ports that belong only to the deleted VLAN are moved to the VLAN identified by the PVID. Tagged ports that belong to multiple VLANs are removed from the deleted VLAN only.

Multiple VLANS with tagging

The following figure shows only those switch port to server links that must be configured for the example. While not shown, all other server links remain set at their default settings.

Figure 7 Multiple VLANs with VLAN tagging

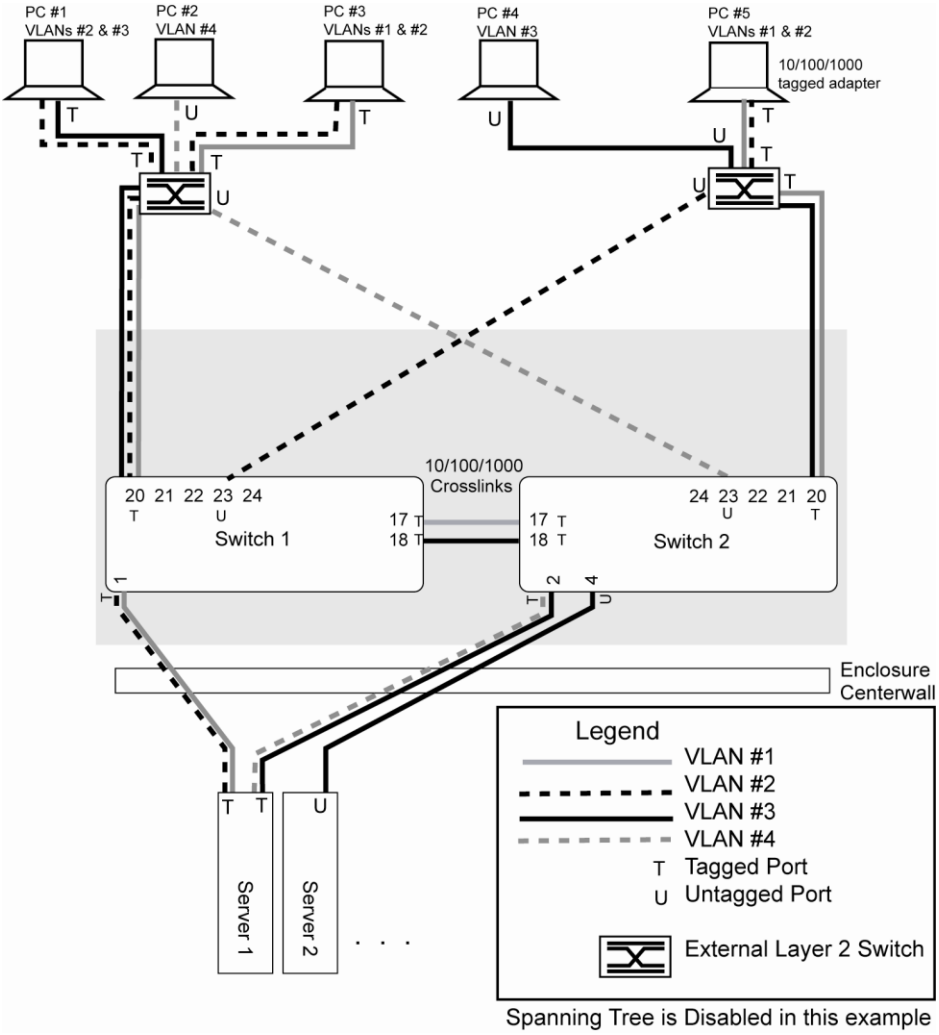

The features of this VLAN are described in the following table:

NOTE: The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your switch model.

Table 9 Multiple VLANs with tagging

Component	Description
Switch 1	Switch 1 is configured for VLANS 1, 2, and 3. Port 1 is tagged to accept traffic from VLANs 1 and 2. Ports 17 and 18 are tagged members of a trunk that accepts traffic from VLANs 1 and 3. Port 20 is tagged to accept traffic from VLANs 1, 2, and 3. Port 23 is an untagged member of VLAN 2.
Switch 2	Switch 2 is configured for VLANS 1, 3, and 4. Port 2 is tagged to accept traffic from VLANS 3 and 4. Port 4 is configured only for VLAN 3, so VLAN tagging is off. Port 20 is tagged to accept traffic from VLANs 1 and 3. Port 23 is an untagged member of VLAN 4.
CPU Blade Server #1	This high-use blade server needs to be accessed from all VLANs and IP subnets. The server has a VLAN-tagging adapter installed with VLAN tagging turned on. One adapter is attached to one of the switch's 10/100/1000 Mbps ports, that is configured for VLANs 1 and 2. One adapter is configured for VLANs 3 and 4. Because of the VLAN tagging capabilities of both the adapter and the switch, the server is able to communicate on all four VLANs in this network while maintaining broadcast separation among all four VLANs and subnets.

Table 9 Multiple VLANs with tagging

Component	Description
CPU Blade Server #2	This blade server belongs to VLAN 3. The port that the VLAN is attached to is configured only for VLAN 3, so VLAN tagging is off.
PC #1	This PC is a member of VLAN 2 and 3. Via VLAN 2, it can communicate with Server 1, PC 3, and PC 5. Via VLAN 3, it can communicate with Server 1, Server 2, and PC 4.
PC #2	This PC is a member of VLAN 4, and can only communicate with Server 1.
PC #3	This PC is a member of VLAN 1 and VLAN 2. Via VLAN 1, it can communicate with Server 1 and PC 5. Via VLAN 2, it can communicate with Server 1, PC 1, and PC 5.
PC #4	This PC is a member of VLAN 3, and it can communicate with Server 1, Server 2, and PC 1.
PC #5	This PC is a member of both VLAN 1 and VLAN 2. Via VLAN 1, it can communicate with Server 1 and PC 3. Via VLAN 2, it can communicate with Server 1, PC 1, and PC 3. The Layer 2 switch port to which it is connected is configured for both VLAN 1 and VLAN 2 and has tagging enabled.

NOTE: All PCs connected to a tagged port must have an Ethernet adapter with VLAN-tagging capability installed.

Configuring the example network

These examples describe how to configure ports and VLANs on Switch 1 and Switch 2.

Configuring ports and VLANs on Switch 1 (AOS CLI example)

To configure ports and VLANs on Switch 1, do the following:

1. On Switch 1, enable VLAN tagging on the necessary ports.

```
Main# /cfg/port 1
>> Port 1# tag e                          (Select port 1: connection to server 1)

Current VLAN tag support: disabled
New VLAN tag support:     enabled          (Enable tagging)
Port 1 changed to tagged.

Main# /cfg/port 17                       (Select crosslink link port 17)
>> Port 17# tag e                        (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 17 changed to tagged.

Main# /cfg/port 18                        (Select crosslink link port 18)
>> Port 18# tag e                         (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 18 changed to tagged.

Main# /cfg/port 20                        (Select uplink port 20)
>> Port 20# tag e                         (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 20 changed to tagged.

>> Port 20# apply                         (Apply the port configurations)
```

2. Configure the VLANs and their member ports. Since all ports are by default configured for VLAN 1, configure only those ports that belong to VLAN 2. crosslink ports 17 and 18 must belong to VLANs 1 and 3.

```
>> /cfg/l2/vlan 2

>> VLAN 2# add 1                                    (Add port 1 to VLAN 2)
Current ports for VLAN 2: empty
Pending new ports for VLAN 2: 1

>> VLAN 2# add 20                                   (Add port 20 to VLAN 2)
Current ports for VLAN 2: 1
Pending new ports for VLAN 2: 20

>> VLAN 2# add 23                                   (Add port 23 to VLAN 2)
Port 23 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 2 [y/n]: y
Current ports for VLAN 2: 1, 20
Pending new ports for VLAN 2: 23

>> /cfg/l2/vlan 3
>> VLAN 3# add 17                                   (Add port 17 to VLAN 3)
Current ports for VLAN 3: empty
Pending new ports for VLAN 3: 17

>> VLAN 3# add 18
Current ports for VLAN 3: 17                        (Add port 18 to VLAN 3)
Pending new ports for VLAN 3: 18

>> VLAN 3# add 20                                   (Add port 20 to VLAN 3)
Current ports for VLAN 3: 17, 18
Pending new ports for VLAN 3: 20

>> /cfg/port 23/tagpvid                             (Disable tagpvid)
Current tag pvid support: enabled
Enter new tag pvid support [d/e]: d
UNTAG on pvid

>> apply                                           (Apply the port configurations)
>> save                                            (Save the port configurations)
```

Configuring ports and VLANs on Switch 2 (AOS CLI example)

To configure ports and VLANs on Switch 2, do the following:

1. On Switch 2, enable VLAN tagging on the necessary ports. Port 4 (connection to server 2) remains untagged, so it is not configured below.

```
Main# /cfg/port 2                          (Select port 2: connection to server 1)
>> Port 2# tag e
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 2 changed to tagged.

Main# /cfg/port 17                         (Select crosslink link port 17)
>> Port 17# tag e                                    (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 17 changed to tagged.

Main# /cfg/port 18                            (Select crosslink link port 18)
>> Port 18# tag e                                      (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 18 changed to tagged.

Main# /cfg/port 20                                   (Select uplink port 20)
>> Port 20# tag e                                           (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 20 changed to tagged.
>> Port 20# apply                             (Apply the port configurations)
```

2. Configure the VLANs and their member ports. Since all ports are by default configured for VLAN 1, configure only those ports that belong to other VLANs.

```
>> /cfg/l2/vlan 3
>> VLAN 3# add 2
Current ports for VLAN 3: empty
Pending new ports for VLAN 3: 2

>> VLAN 3# add 4
Current ports for VLAN 3: 2
Pending new ports for VLAN 3: 17

>> VLAN 3# add 17
Current ports for VLAN 3: 2, 4
Pending new ports for VLAN 3: 17

>> VLAN 3# add 18
Current ports for VLAN 3: 2, 17
Pending new ports for VLAN 3: 18

>> VLAN 3# add 20
Current ports for VLAN 3: 2, 17, 18
Pending new ports for VLAN 3: 20


>> /cfg/l2/vlan 4
>> VLAN 4# add 2
Current ports for VLAN 4: empty
Pending new ports for VLAN 4: 2

>> VLAN 4# add 23
Port 23 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 4 [y/n]: y
Current ports for VLAN 4: 2
Pending new ports for VLAN 4: 23

>> /cfg/port 4/tagpvid
Current tag pvid support: enabled
Enter new tag pvid support [d/e]: d
UNTAG on pvid

>> /cfg/port 23/tagpvid
Current tag pvid support: enabled
Enter new tag pvid support [d/e]: d
UNTAG on pvid


>> apply                                    (Apply the port configurations)
>> save                                     (Save the port configurations)
```

The external Layer 2 switches should also be configured for VLANs and tagging.

Configuring ports and VLANs on Switch 1 (BBI example)

To configure ports and VLANs on Switch 1, do the following:

1. On the switch 1, enable VLAN tagging on the necessary ports.
   a. Click the Configure context button on the Toolbar.
   b. Open the Switch folder, and select Switch Ports (click the underlined text, not the folder).

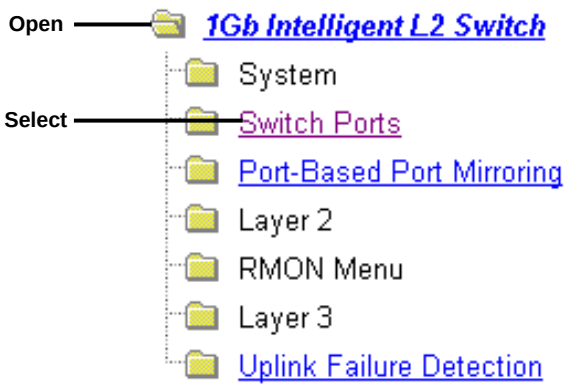

c. Click a port number to select it.

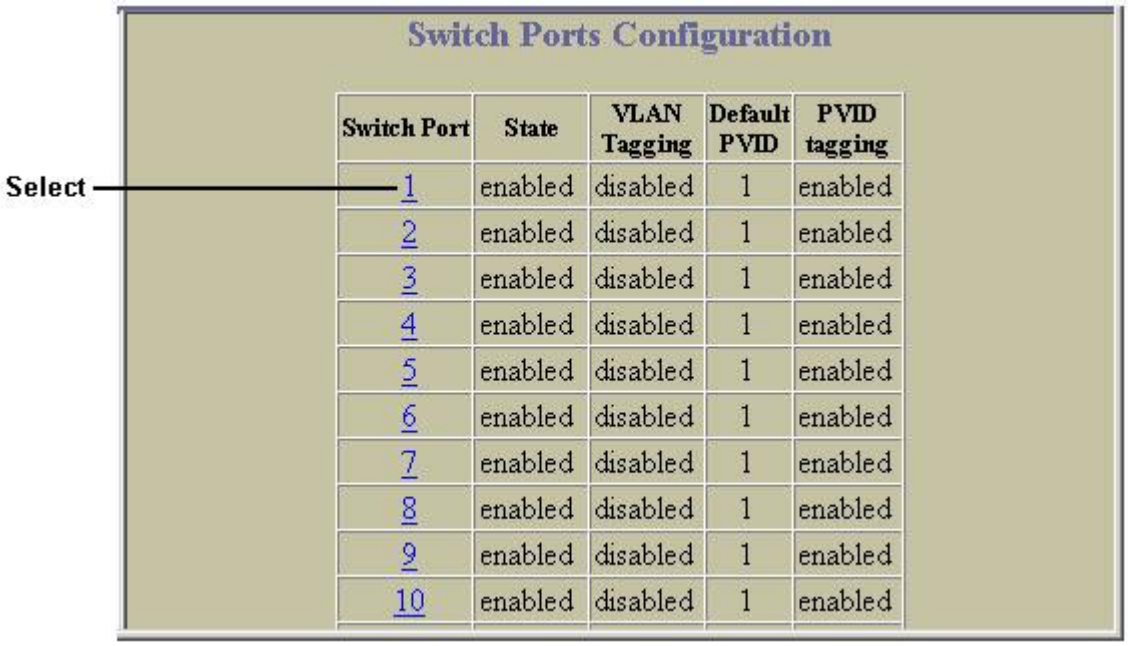

d. Enable the port and enable VLAN tagging.

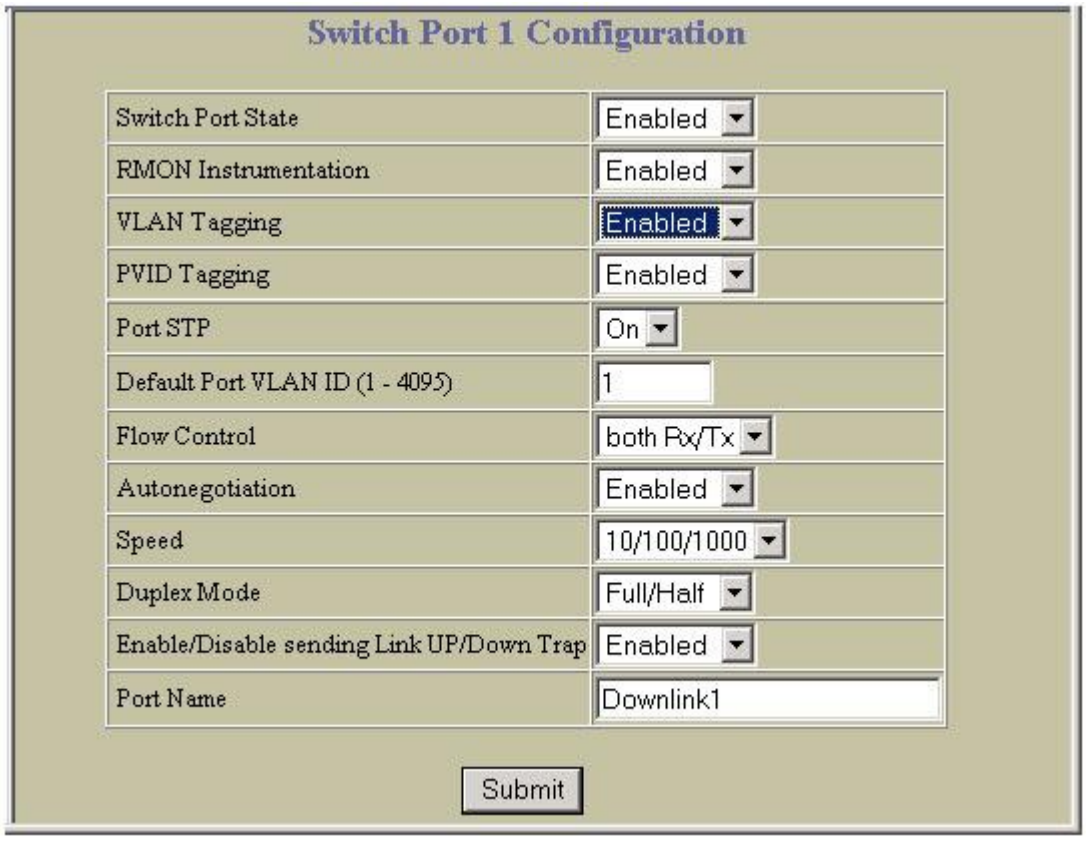

e. Click Submit.

2. Configure the VLANs and their member ports.
   a. Open the Virtual LANs folder, and select Add VLAN.

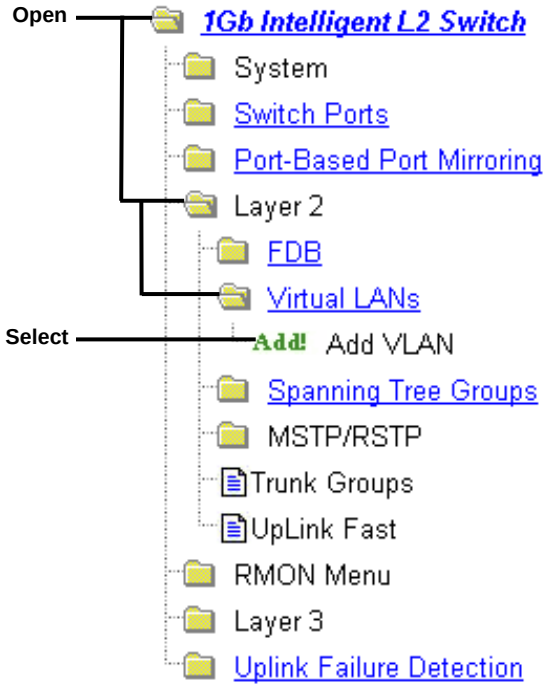

   b. Enter the VLAN name, VLAN ID number, and enable the VLAN. To add ports, select each port in the Ports Available list and click Add. Since all ports are configured for VLAN 1 by default, configure only those ports that belong to VLAN 2. The crosslink ports 17 and 18 must belong to VLANs 1 and 2.

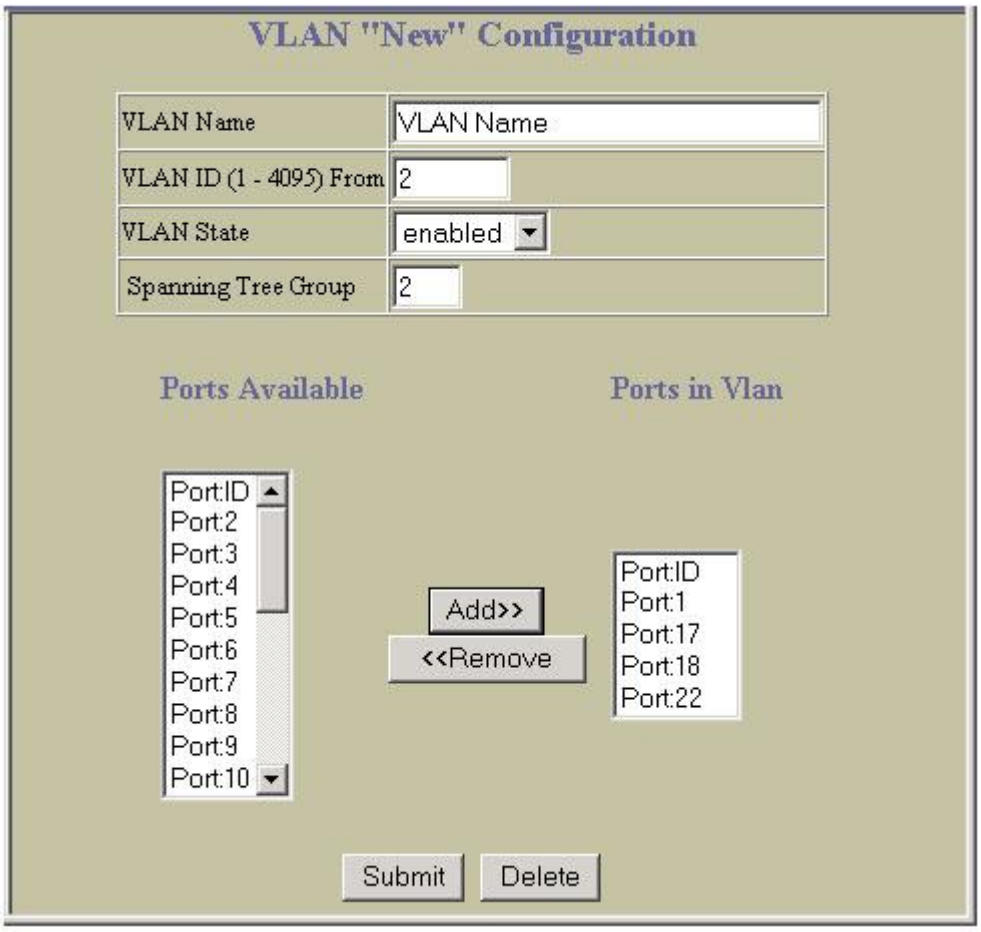

   c. Click Submit.

The external Layer 2 switches should also be configured for VLANs and tagging.

3. Apply, verify, and save the configuration.

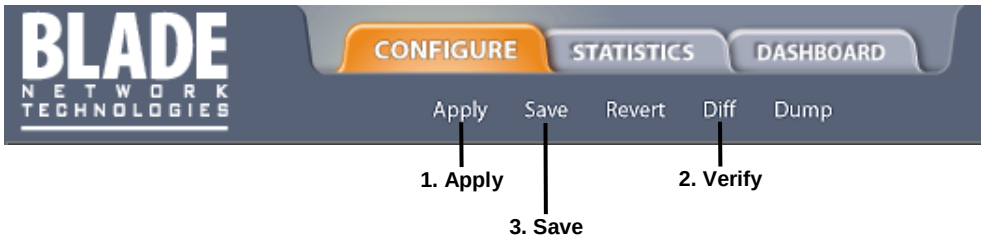

FDB static entries

Static entries in the Forwarding Database (FDB) allow the switch to forward packets without flooding ports to perform a lookup. A FDB static entry is a MAC address associated with a specific port and VLAN. The switch supports 128 static entries. Static entries are manually configured, using the `/cfg/l2/fdb/static` command.

FDB static entries are permanent, so the FDB Aging value does not apply to them. Static entries are manually added to the FDB, and manually deleted from the FDB.

Incoming frames that contain the static entry as the source MAC can use only ports configured for the static entry.

Trunking support for FDB static entries

A FDB static entry can be added to a port that is a member of a trunk group, as follows:

- Static (manually configured) trunk group
- Dynamic (LACP) trunk group

The trunk group supports the FDB static entry. If the port with the static entry fails, other ports in the trunk handle the traffic. If the port is removed from the trunk, the static entry is removed from the trunk, but remains configured on the port.

The FDB information commands (`/info/l2/fdb`) display trunk support for static FDB entries, if applicable:

```
>> Forwarding Database# dump
     MAC address      VLAN  Port  Trnk  State
  -----------------   ----  ----  ----  -----
  00:00:2e:9b:db:f8      1           1   TRK
  00:00:5e:00:01:f4      1  24           FWD
  00:01:81:2e:b5:60      1  24           FWD
  00:02:a5:e9:76:30      1           1   TRK
  00:03:4b:e2:15:f1      1  24           FWD
```

Configuring a static FDB entry

Perform the following actions to configure a static FDB entry:

```
Main# /cfg/l2/fdb/static                          (Select static FDB menu)
>> Static FDB# add 00:60:af:00:02:30
Enter VLAN number: 2
Enter port (1-24): 2
>> Static FDB# apply                              (Apply the configuration)
>> Static FDB# save                               (Save the configuration)
```

Spanning Tree Protocol

Introduction

When multiple paths exist on a network, Spanning Tree Protocol (STP) configures the network so that a switch uses only the most efficient path. The following topics are discussed in this chapter:

- Overview
- Bridge Protocol Data Units (BPDUs)
- Spanning Tree Group (STG) configuration guidelines
- Multiple Spanning Trees

Overview

Spanning Tree Protocol (STP) detects and eliminates logical loops in a bridged or switched network. STP forces redundant data paths into a standby (blocked) state. When multiple paths exist, STP configures the network so that a switch uses only the most efficient path. If that path fails, STP automatically sets up another active path on the network to sustain network operations.

The switch supports IEEE 802.1D Spanning Tree Protocol for STG 1, and Per VLAN Spanning Tree Protocol (PVST+) for STGs 2-32, by default.

NOTE: The switch also supports IEEE 802.1w Rapid Spanning Tree Protocol and IEEE 802.1s Multiple Spanning Tree Protocol. For more information, see the “RSTP and MSTP” chapter in this guide.

Bridge Protocol Data Units

To create a spanning tree, the application switch generates a configuration Bridge Protocol Data Unit (BPDU), which it then forwards out of its ports. All switches in the Layer 2 network participating in the spanning tree gather information about other switches in the network through an exchange of BPDUs.

A BPDU is a 64-byte packet that is sent out at a configurable interval, which is typically set for two seconds. The BPDU is used to establish a path, much like a “hello” packet in IP routing. BPDUs contain information about the transmitting bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and port path cost. If the ports are tagged, each port sends out a special BPDU containing the tagged information.

The generic action of a switch on receiving a BPDU is to compare the received BPDU to its own BPDU that it will transmit. If the received BPDU has a priority value closer to zero than its own BPDU, it will replace its BPDU with the received BPDU. Then, the application switch adds its own bridge ID number and increments the path cost of the BPDU. The application switch uses this information to block any redundant paths.

Determining the path for forwarding BPDUs

When determining which port to use for forwarding and which port to block, the switch uses information in the BPDU, including each bridge priority ID. A technique based on the “lowest root cost” is then computed to determine the most efficient path for forwarding.

Bridge priority

The bridge priority parameter controls which bridge on the network is the STP root bridge. To make one switch the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. The bridge priority is configured using the `/cfg/l2/stp/brg/prior` command in the CLI.

Port priority

The port priority helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. The port priority is configured using the `/cfg/l2/stp/port x/prior` command in the CLI.

Port path cost

The port path cost assigns lower values to high-bandwidth ports, such as Gigabit Ethernet, to encourage their use. The objective is to use the fastest links so that the route with the lowest cost is chosen. A value of 0 indicates that port cost is computed dynamically based on link speed. This works when forcing link speed, so it does not just apply to “auto negotiated link speed”.

By default, all switch ports have the path cost set to 4, independent of the link speed. To use dynamic path cost, based on link speed, set the path cost to 0 (zero). For example, if the path cost is set to zero:

- A 100 Mbps link receives a path cost of 19
- A 10 Mbps link receives a path cost of 100

Spanning Tree Group configuration guidelines

This section provides important information on configuring Spanning Tree Groups (STGs).

Default Spanning Tree configuration

In the default configuration, a single STG with the ID of 1 includes all ports except Port 19 on the switch. It is called the default STG. All other STGs (except the default STG) are empty, and VLANs must be added by the user.

You cannot assign ports directly to an STG. Add the ports to a VLAN, and add the VLAN to the STG. STGs 1-31 are enabled by default and assigned an ID number from 1 to 31. STG 32 is disabled by default, and contains the management VLAN 4095.

An STG cannot be deleted, only disabled. If you disable the STG while it still contains VLAN members, Spanning Tree will be off on all ports belonging to that VLAN.

Adding a VLAN to a Spanning Tree Group

If no VLANs exist beyond the default VLAN 1, see the “Creating a VLAN” section in this chapter for information on adding ports to VLANs.

Add the VLAN to the STG using the command `/cfg/l2/stp <stg number>/add <vlan number>`.

Creating a VLAN

When you create a VLAN, then that VLAN automatically belongs to STG 1, the default STG. If you want the VLAN in another STG, you must move the VLAN by assigning it to another STG.

To move a newly created VLAN to an existing STG:

1. Create the VLAN.
2. Add the VLAN to an existing STG.

When creating a VLAN also consider the following:

- A VLAN cannot belong to more than one STG.
- VLANs that span multiple switches must be mapped within the same Spanning Tree Group (have the same STG ID) across all the switches.

Rules for VLAN tagged ports

Rules for VLAN tagged ports are listed below:

- If a port is tagged, it can belong to multiple STGs.
- When a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.
- An untagged port cannot span multiple STGs.

Adding and removing ports from STGs

Information on adding and removing ports from STGs is as follows:

- By default, all ports except Port 19 belong to VLAN 1 and STG 1.
- Each port is always a member of at least one VLAN. Each VLAN is always a member of at least one STG. Port membership within VLANs can be changed, and VLAN membership within STGs can be changed. To move a port from one STG to another, move the VLAN to which the port belongs, or move the port to a VLAN that belongs to the STG.
- When you remove a port from a VLAN, that port is also removed from the STG to which the VLAN belongs. However, if that port belongs to another VLAN in the same STG, the port remains in the STG.
- If you remove an untagged port from a non-default VLAN and STG, it is added to VLAN 1 and STG 1.

The relationship between ports, trunk groups, VLANs, and spanning trees is shown in the following table.

Table 10 Ports, trunk groups, and VLANs

Switch element	Belongs to
Port	Trunk group, or one or more VLANs
Trunk group	One or more VLANs
VLAN (non-default)	One Spanning Tree Group

Assigning cost to ports and trunk groups

When you configure a trunk group to participate in a Spanning Tree Group, manually assign port and trunk costs to ensure that each trunk group has a lower STP cost than the cost of each port within the trunk group. This ensures that the trunk group remains in the Forwarding state.

Multiple Spanning Trees

Each switch supports a maximum of 32 Spanning Tree Groups (STGs). Multiple STGs provide multiple data paths, which can be used for load-balancing and redundancy.

You enable independent links on two switches using multiple STGs by configuring each path with a different VLAN and then assigning each VLAN to a separate STG. Each STG is independent. Each STG sends its own Bridge Protocol Data Units (BPDUs), and each STG must be independently configured.

The STG, or bridge group, forms a loop-free topology that includes one or more virtual LANs (VLANs). The switch supports 32 STGs running simultaneously. The default STG 1 supports IEEE 802.1d Spanning Tree Protocol, and may contain more than one VLAN. All other STGs support Per VLAN Spanning Tree (PVST+), and may contain only one VLAN each. The switch can support multiple VLANs in STGs 2-32; however, you must enable IEEE 802.1s Multiple Spanning Tree Protocol mode. For more information, see the “RSTP and MSTP” chapter in this guide.

Why do we need Multiple Spanning Trees?

The following figure shows a simple example of why we need multiple Spanning Trees. This example assumes that port 17 and 18 are not part of Trunk Group 1. Two VLANs (VLAN 1 and VLAN 2) exist between Switch 1 and Switch 2. If the same Spanning Tree Group is enabled on both switches, the switches see an apparent loop and block port 18 on Switch 2, which cuts off communication between the switches for VLAN 2.

Figure 8 Two VLANs on one instance of Spanning Tree Protocol

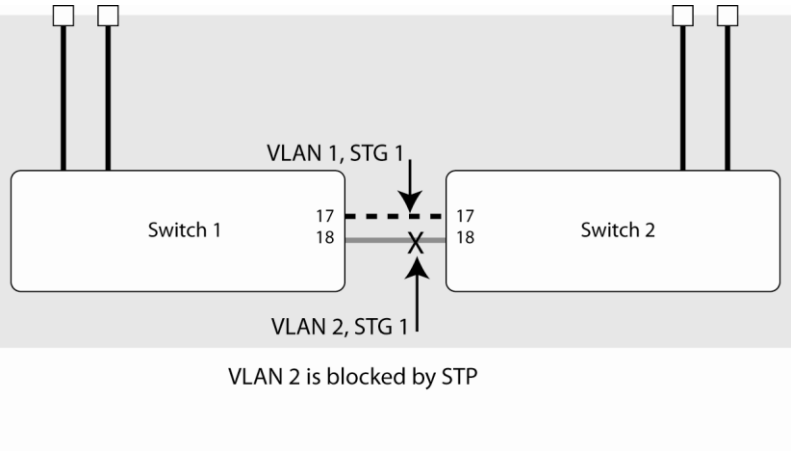

In the following figure, VLAN 1 and VLAN 2 belong to different Spanning Tree Groups. The two instances of spanning tree separate the topology without forming a loop, so that both VLANs can forward packets between the switches without losing connectivity.

Figure 9 Two VLANs on separate instances of Spanning Tree Protocol

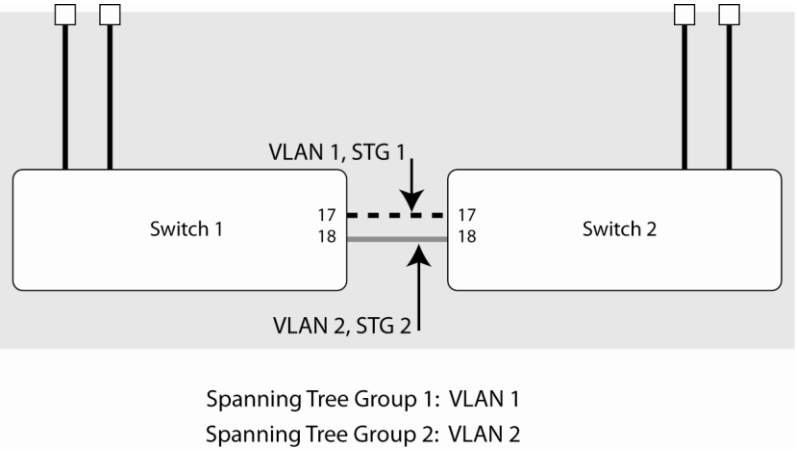

VLAN participation in Spanning Tree Groups

The following table shows which switch ports participate in each Spanning Tree Group. By default, server ports (ports 1-16) do not participate in Spanning Tree, even though they are members of their respective VLANs.

Table 11 VLAN participation in Spanning Tree Groups

	VLAN 1	VLAN 2
Switch 1	Spanning Tree Group 1 Port 17	Spanning Tree Group 2 Port 18
Switch 2	Spanning Tree Group 1 Port 17	Spanning Tree Group 2 Port 18

Configuring Multiple Spanning Tree Groups

This section explains how to assign each VLAN to its own Spanning Tree Group on the switches 1 and 2.

By default, Spanning Tree Groups 2-31 are empty, and Spanning Tree Group 1 contains all configured VLANs (except VLAN 4095) until individual VLANs are explicitly assigned to other Spanning Tree Groups. Except for the default Spanning Tree Group 1, which may contain more than one VLAN, Spanning Tree Groups 2-32 may contain only one VLAN each.

NOTE: Each instance of Spanning Tree Group is enabled by default.

Configuring Switch 1 (AOS CLI example)

1. Configure port and VLAN membership on Switch 1 as described in the “Configuring ports and VLANs on Switch 1 (AOS CLI example)” section, in the “VLANs” chapter of this guide.
2. Add VLAN 2 to Spanning Tree Group 2.

```
>> /cfg/l2/stp 2                          (Select Spanning Tree Group 2)
>> Spanning Tree Group 2# add 2                 (Add VLAN 2)
```

   VLAN 2 is automatically removed from spanning tree group 1.
3. Apply and save.

```
>> apply                                 (Apply the port configurations)
>> save                                  (Save the port configurations)
```

Configuring Switch 2 (AOS CLI example)

1. Configure port and VLAN membership as described in the “Configuring ports and VLANs on Switch 2 (CLI example)” section in the “VLANs” chapter of this guide.
2. Add VLAN 2 to Spanning Tree Group 2.

```
>> /cfg/l2/stp 2                          (Select Spanning Tree Group 2)
>> Spanning Tree Group 2# add 2           (Add VLAN 2)
```

3. VLAN 2 is automatically removed from Spanning Tree Group 1.
4. Apply and save.

```
>> apply                                 (Apply the port configurations)
>> save                                  (Save the port configurations)
```

Configuring Switch 1 (BBI example)

1. Configure port and VLAN membership on Switch 1 as described in the “Configuring ports and VLANs on Switch 1 (BBI example)” section, in the “VLANs” chapter of this guide.
2. Add VLAN 2 to Spanning Tree Group 2.
   a. Click the Configure context button on the Toolbar.
   b. Open the Spanning Tree Groups folder, and select Add Spanning Tree Group.

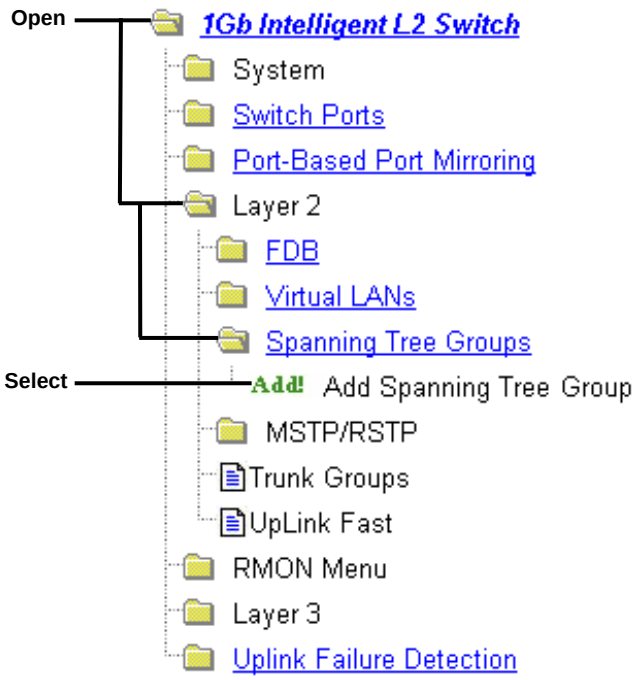

c. Enter the Spanning Tree Group number and set the Switch Spanning Tree State to on. To add a VLAN to the Spanning Tree Group, select the VLAN in the VLANs Available list, and click Add. VLAN 2 is automatically removed from Spanning Tree Group 1.

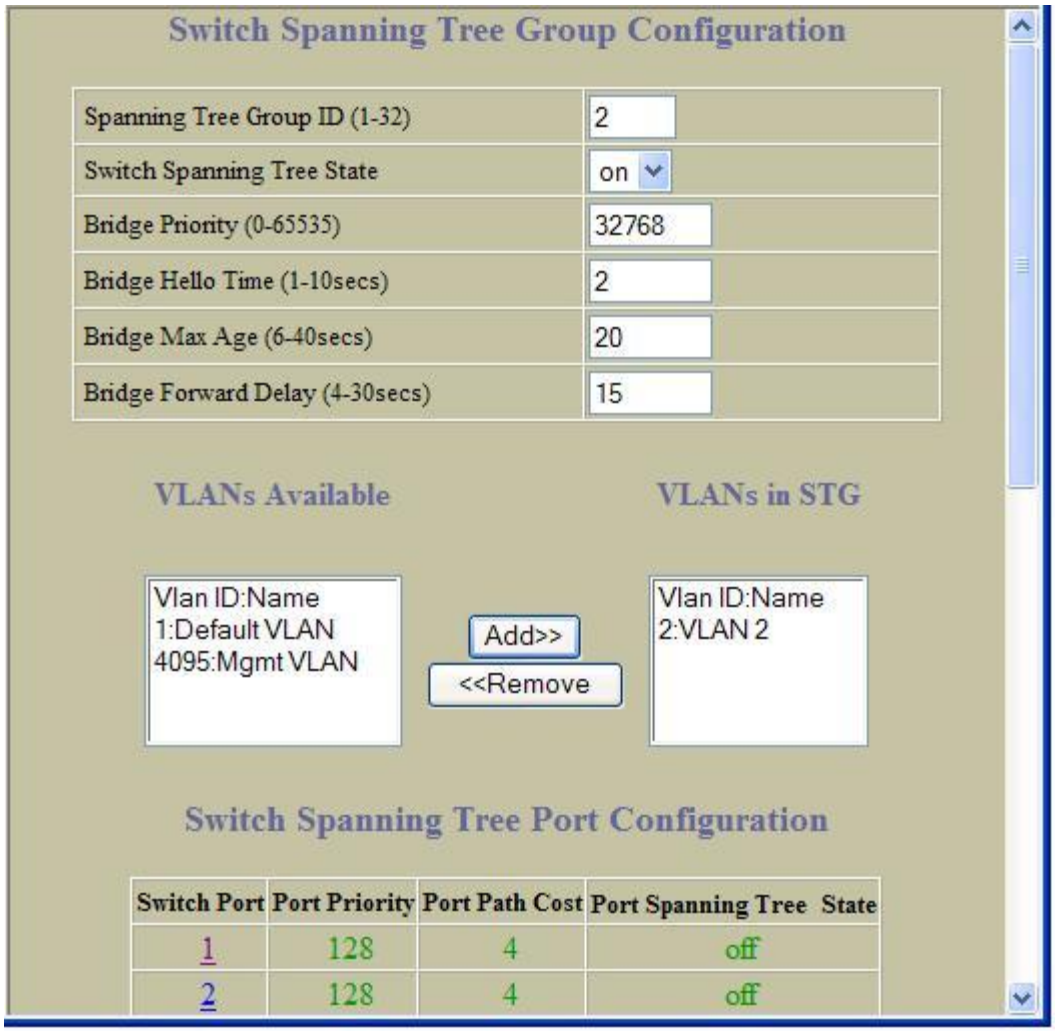

d. Scroll down, and click Submit.

3. Apply, verify, and save the configuration.

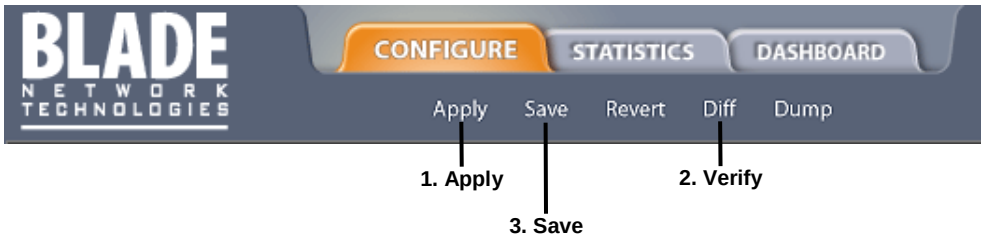

Port Fast Forwarding

Port Fast Forwarding permits a port that participates in Spanning Tree to bypass the Listening and Learning states and enter directly into the Forwarding state. While in the Forwarding state, the port listens to the BPDUs to learn if there is a loop and, if dictated by normal STG behavior (following priorities, etc.), the port transitions into the Blocking state.

This feature permits the switch to interoperate well with Fast Path, a NIC Teaming feature.

Configuring Port Fast Forwarding

Use the following CLI commands to enable Port Fast Forwarding on an external port.

```
>> # /cfg/l2/stp 1/port 20            (Select port 20)
>> Spanning Tree Port 20# fastfwd ena (Enable Port Fast Forwarding)
>> Spanning Tree Port 20# apply       (Make your changes active)
>> Spanning Tree Port 20# save        (Save for restore after reboot)
```

Fast Uplink Convergence

Fast Uplink Convergence enables the switch to quickly recover from the failure of the primary link or trunk group in a Layer 2 network using Spanning Tree Protocol. Normal recovery can take as long as 60 seconds, while the backup link transitions from Blocking to Listening to Learning and then Forwarding states. With Fast Uplink Convergence enabled, the switch immediately places the secondary path into Forwarding state, and sends multicasts of addresses in the forwarding database (FDB) and ARP table over the secondary link so that upstream switches can learn the new path.

Configuration guidelines

When you enable Fast Uplink Convergence, the switch software automatically makes the following configuration changes:

- Increases the bridge priority to 65500 so that it does not become the root switch.
- Increases the cost of all of the external ports by 3000, across all VLANs and Spanning Tree Groups. This ensures that traffic never flows through the switch to get to another switch unless there is no other path.

When you disable Fast Uplink Convergence, the bridge priorities and path cost are set to their default values for all STP groups.

Configuring Fast Uplink Convergence

Use the following CLI commands to enable Fast Uplink Convergence on external ports:

```
>> # /cfg/l2/upfast ena   (Enable Fast Uplink convergence)
>> Layer 2# apply         (Make your changes active)
>> Layer 2# save          (Save for restore after reboot)
```

RSTP and MSTP

Introduction

Rapid Spanning Tree Protocol (IEEE 802.1w) enhances the Spanning Tree Protocol (IEEE 802.1D) to provide rapid convergence on Spanning Tree Group 1. Multiple Spanning Tree Protocol (IEEE 802.1s) extends the Rapid Spanning Tree Protocol to provide both rapid convergence and load balancing in a VLAN environment.

The following topics are discussed in this chapter:

- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)

Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) provides rapid convergence of the spanning tree and provides for fast reconfiguration critical for networks carrying delay-sensitive traffic such as voice and video. RSTP significantly reduces the time to reconfigure the active topology of the network when changes occur to the physical topology or its configuration parameters. RSTP reduces the bridged-LAN topology to a single Spanning Tree.

For more information about Spanning Tree Protocol, see the “Spanning Tree Protocol” chapter in this guide.

RSTP parameters are configured in Spanning Tree Group 1. STP Groups 2-32 do not apply to RSTP, and must be cleared. There are new STP parameters to support RSTP, and some values to existing parameters are different.

RSTP is compatible with devices that run 802.1D Spanning Tree Protocol. If the switch detects 802.1D BPDUs, it responds with 802.1D-compatible data units. RSTP is not compatible with Per VLAN Spanning Tree (PVST) protocol.

Port state changes

The port state controls the forwarding and learning processes of Spanning Tree. In RSTP, the port state has been consolidated to the following: discarding, learning, and forwarding.

Table 12 RSTP vs. STP port states

Port operational status	STP port state	RSTP port state
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding
Disabled	Disabled	Discarding

Port type and link type

Spanning Tree Configuration includes the following parameters to support RSTP and MSTP:

- Edge port
- Link type

Although these parameters are configured for Spanning Tree Groups 1-32 (`/cfg/l2/stp x/port x`), they only take effect when RSTP/MSTP is turned on.

Edge port

A port that connects to a server or stub network is called an edge port. Therefore, ports 1-16 should have edge enabled. (The default for Ports 1-16 is enabled.) Edge ports can start forwarding as soon as the link is up.

Edge ports do not take part in Spanning Tree, and should not receive BPDUs. If a port with edge enabled does receive a BPDU, it begins STP processing until it is re-enabled.

Link type

The link type determines how the port behaves in regard to Rapid Spanning Tree. The link type corresponds to the duplex mode of the port. A full-duplex link is point-to-point (p2p), while a half-duplex link should be configured as shared. If you select auto as the link type, the port dynamically configures the link type.

RSTP configuration guidelines

This section provides important information about configuring Rapid Spanning Tree Groups:

- When RSTP is turned on, STP parameters apply only to STP Group 1.
- When RSTP is turned on, all VLANs from STP Groups other than STP Group 1 are moved to STP Group 1. The other STP Groups (2-32) are turned off.

RSTP configuration example

This section provides steps to configure Rapid Spanning Tree on the switch, using the Command Line Interface (CLI) or the Browser-based Interface (BBI).

Configuring Rapid Spanning Tree (CLI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (CLI example)” section in the “VLANs” chapter of this guide.
2. Set the Spanning Tree mode to Rapid Spanning Tree.

```
>> /cfg/l2/mrst                      (Select Multiple Spanning Tree menu)
>> Multiple Spanning Tree# mode rstp   (Set mode to Rapid Spanning Tree)
>> Multiple Spanning Tree# on              (Turn Rapid Spanning Tree on)
```

3. Apply and save the changes.

```
>> # apply               (Apply the configuration)
>> # save                (Save the configuration)
```

Configuring Rapid Spanning Tree Protocol (BBI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (BBI example)” section in the “VLANs” chapter of this guide.
2. Configure RSTP general parameters.
   a. Click the Configure context button on the Toolbar.

b. Open the MSTP/RSTP folder, and select General.

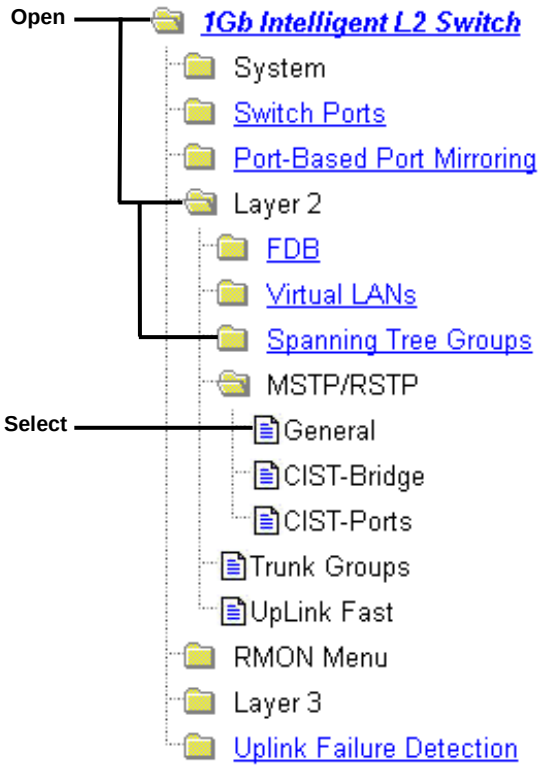

c. Select RSTP mode, and set the MSTP/RSTP state to ON.

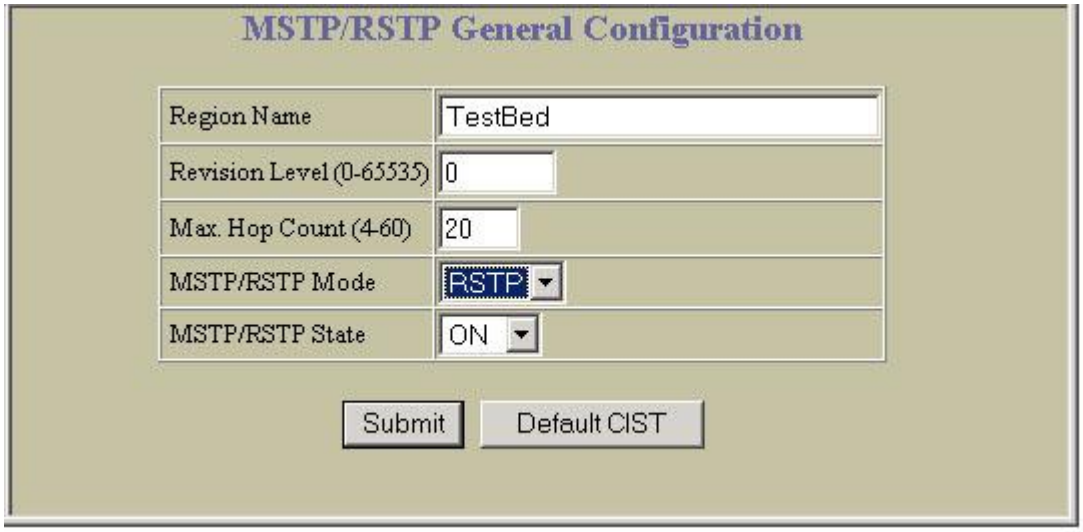

d. Click Submit.

3. Apply, verify, and save the configuration.

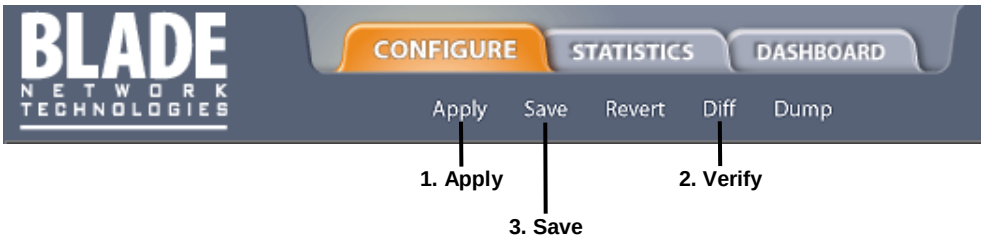

Multiple Spanning Tree Protocol

IEEE 802.1s Multiple Spanning Tree extends the IEEE 802.1w Rapid Spanning Tree Protocol through multiple Spanning Tree Groups. MSTP maintains up to 32 spanning-tree instances that correspond to STP Groups 1-32.

In Multiple Spanning Tree Protocol (MSTP), several VLANs can be mapped to each Spanning-Tree instance. Each Spanning-Tree instance is independent of other instances. MSTP allows frames assigned to different VLANs to follow separate paths, each path based on an independent Spanning-Tree instance. This approach provides multiple forwarding paths for data traffic, enabling load balancing, and reducing the number of Spanning-Tree instances required to support a large number of VLANs.

MSTP region

A group of interconnected bridges that share the same attributes is called an MSTP region. Each bridge within the region must share the following attributes:

- Alphanumeric name
- Revision level
- VLAN-to-STG mapping scheme

MSTP provides rapid reconfiguration, scalability, and control due to the support of regions, and multiple Spanning-Tree instances support within each region.

Common Internal Spanning Tree

The Common Internal Spanning Tree (CIST) provides a common form of Spanning Tree Protocol, with one Spanning Tree instance that can be used throughout the MSTP region. CIST allows the switch to interoperate with legacy equipment, including devices that run IEEE 802.1D (STP).

CIST allows the MSTP region to act as a virtual bridge to other bridges outside of the region, and provides a single Spanning-Tree instance to interact with them.

CIST is the default spanning tree group. When VLANs are removed from STG 1-32, the VLANs automatically become members of the CIST.

CIST port configuration includes Hello time, Edge port status (enable/disable), and Link Type. These parameters do not affect Spanning Tree Groups 1-32. They apply only when the CIST is used.

MSTP configuration guidelines

This section provides important information about configuring Multiple Spanning Tree Groups:

- When you turn on MSTP, the switch automatically moves VLAN 1 to the Common Internal Spanning Tree (CIST).
- Region Name and revision level must be configured. Each bridge in the region must have the same name and revision level.
- The VLAN and STP Group mapping must be the same across all bridges in the region.
- You can move any VLAN to the CIST.
- You can move VLAN 1 into any Spanning Tree Group.

MSTP configuration example

This section provides steps to configure Multiple Spanning Tree Protocol on the switch, using the Command Line Interface (CLI) or the Browser-based Interface (BBI).

Configuring Multiple Spanning Tree Protocol (CLI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (CLI example)” section in the “VLANs” chapter of this guide.
2. Set the mode to Multiple Spanning Tree, and configure MSTP region parameters.

```
>> /cfg/l2/ mrst                     (Select Multiple Spanning Tree menu)
>> Multiple Spanning Tree# mode mstp  (Set mode to
                                       Multiple Spanning Trees)
>> Multiple Spanning Tree# on         (Turn Multiple Spanning Trees on)
>> Multiple Spanning Tree# name xxxxxx (Define the Region name)
>> Multiple Spanning Tree: rev xx     (Define the Region revision level)
```

3. Assign VLANs to Spanning Tree Groups.

```
>> /cfg/l2/stp 2                        (Select Spanning Tree Group 2)
>> Spanning Tree Group 2# add 2         (Add VLAN 2)
>> Spanning Tree Group 2# apply         (Apply the configurations)
```

Configuring Multiple Spanning Tree Protocol (BBI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (BBI example)” section in the “VLANs” chapter of this guide.
2. Configure MSTP general parameters.
   a. Click the Configure context button on the Toolbar.
   b. Open the MSTP/RSTP folder, and select General.

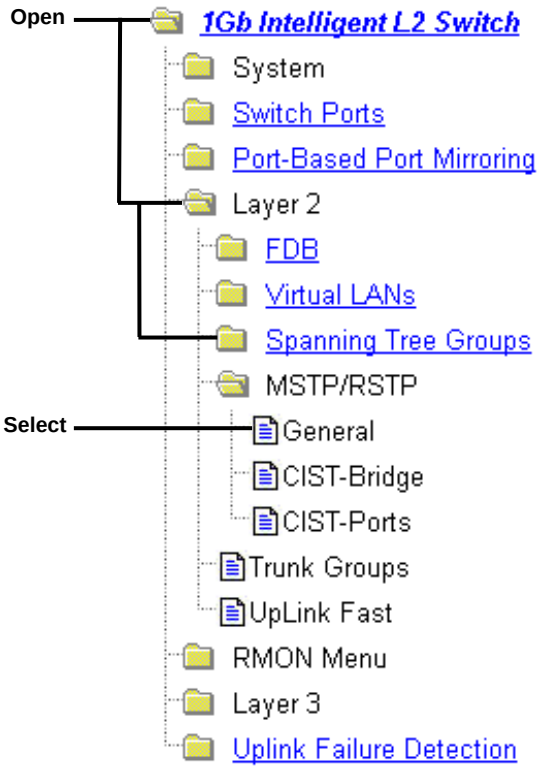

   c. Enter the region name and revision level. Select MSTP mode, and set the MSTP/RSTP state to ON.

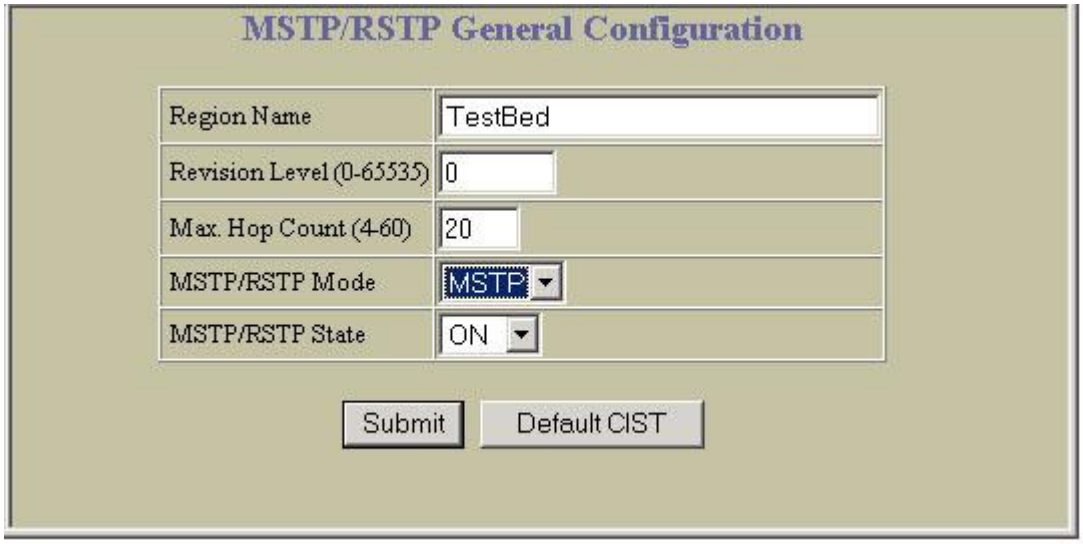

   d. Click Submit.

3. Configure Common Internal Spanning Trees (CIST) bridge parameters.
   a. Open the MSTP/RSTP folder, and select CIST-Bridge.

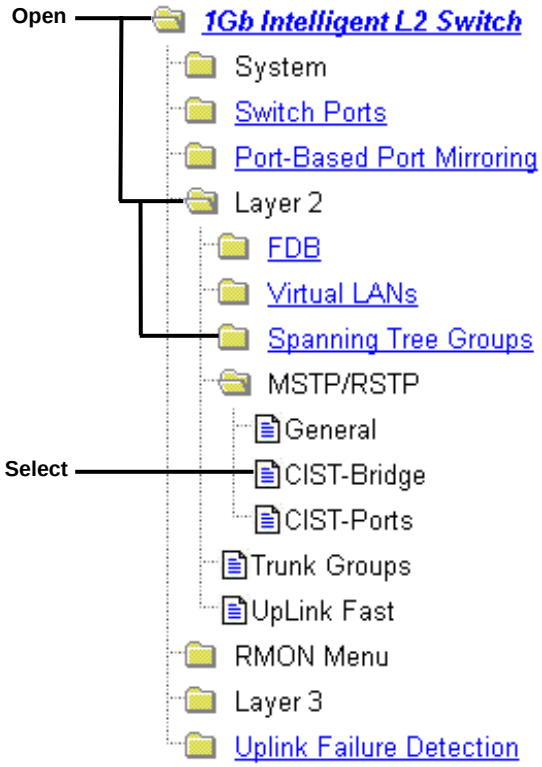

   b. Enter the Bridge Priority, Maximum Age, and Forward Delay values.

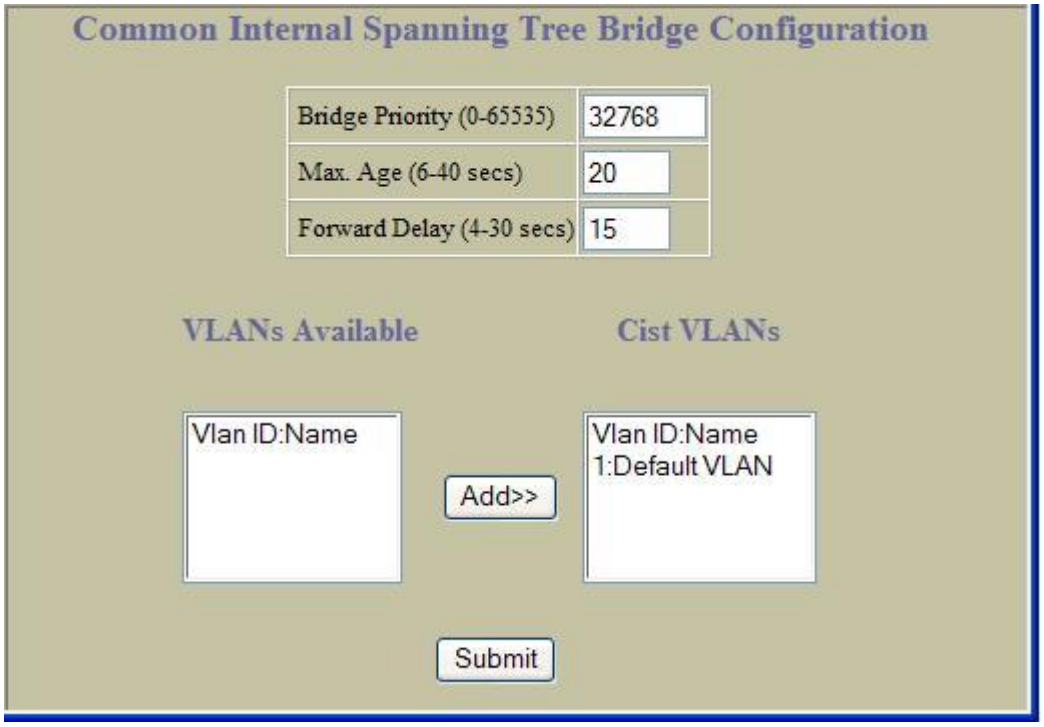

   c. Click Submit.

4. Configure Common Internal Spanning Tree (CIST) port parameters.

   a. Open the MSTP/RSTP folder, and select CIST-Ports.

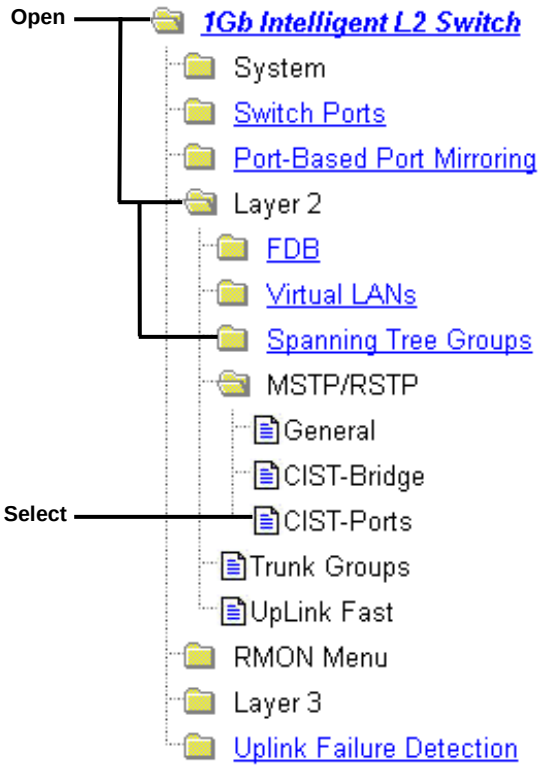

   b. Click a port number to select it.

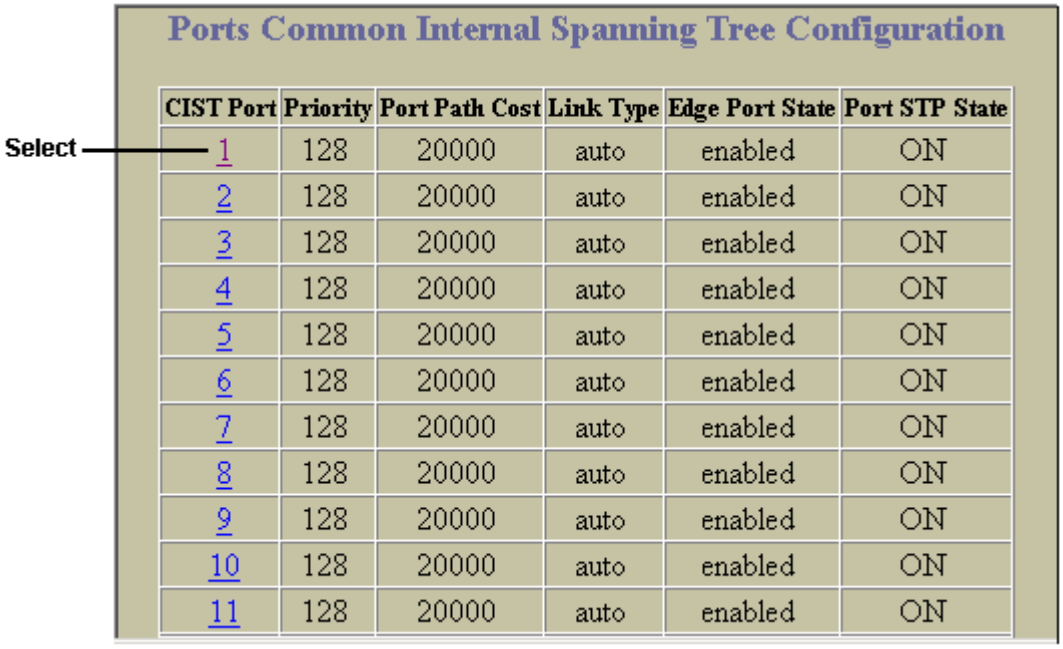

CIST Port	Priority	Port Path Cost	Link Type	Edge Port State	Port STP State
1	128	20000	auto	enabled	ON
2	128	20000	auto	enabled	ON
3	128	20000	auto	enabled	ON
4	128	20000	auto	enabled	ON
5	128	20000	auto	enabled	ON
6	128	20000	auto	enabled	ON
7	128	20000	auto	enabled	ON
8	128	20000	auto	enabled	ON
9	128	20000	auto	enabled	ON
10	128	20000	auto	enabled	ON
11	128	20000	auto	enabled	ON

c. Enter the Port Priority, Path Cost, and select the Link Type. Set the CIST Port State to ON.

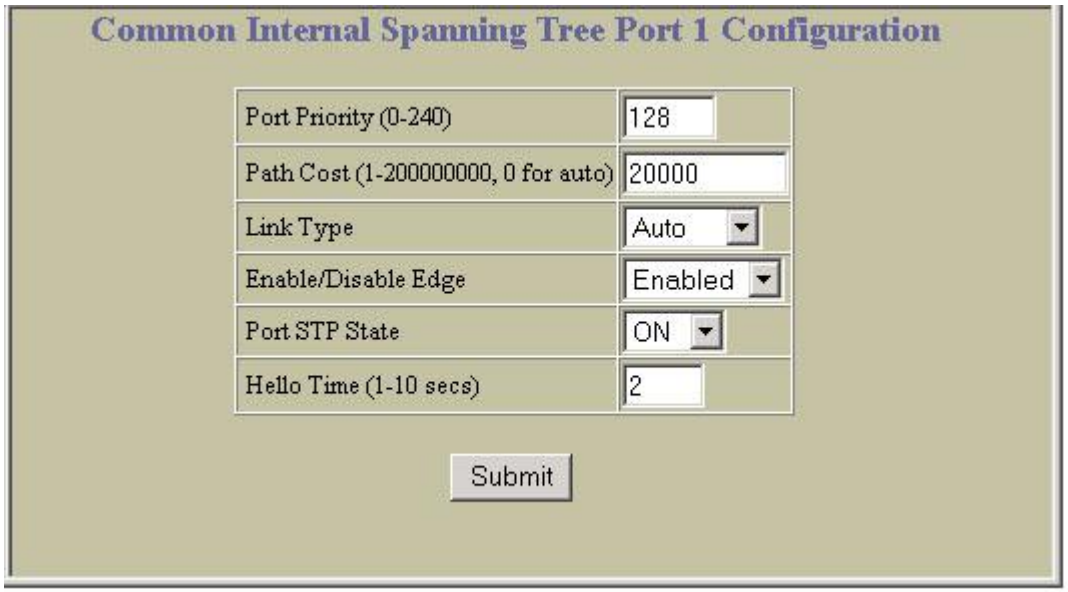

d. Click Submit.

5. Apply, verify, and save the configuration.

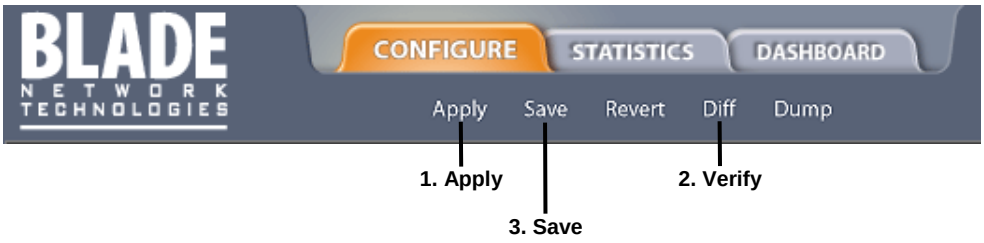

IGMP Snooping

Introduction

IGMP Snooping allows the switch to forward multicast traffic only to those ports that request it. IGMP Snooping prevents multicast traffic from being flooded to all data ports. The switch learns which server hosts are interested in receiving multicast traffic, and forwards it only to ports connected to those servers.

The following topics are discussed in this chapter:

- Overview
- FastLeave
- IGMP Filtering
- Static Multicast Router
- IGMP Snooping Configuration example

Overview

Internet Group Management Protocol (IGMP) is used by IP Multicast routers to learn about the existence of host group members on their directly attached subnet (see RFC 2236). The IP Multicast routers get this information by broadcasting IGMP Query Reports and listening for IP hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IP Multicast source that provides the data streams and the clients that want to receive the data.

IGMP Snooping conserves bandwidth. With IGMP Snooping, the switch learns which ports are interested in receiving multicast data, and forwards multicast data only to those ports. In this way, other ports are not burdened with unwanted multicast traffic.

The switch currently supports snooping for IGMP version 1 and version 2.

The switch can sense IGMP Membership Reports from attached host servers and act as a proxy to set up a dedicated path between the requesting host and a local IP Multicast router. After the pathway is established, the switch blocks the IP Multicast stream from flowing through any port that does not connect to a host member, thus conserving bandwidth.

The client-server path is set up as follows:

- An IP Multicast Router (Mrouter) sends Membership Queries to the switch, which forwards them to all ports in a given VLAN.
- Hosts that want to receive the multicast data stream send Membership Reports to the switch, which sends a proxy Membership Report to the Mrouter.
- The switch sets up a path between the Mrouter and the host, and blocks all other ports from receiving the multicast.
- Periodically, the Mrouter sends Membership Queries to ensure that the host wants to continue receiving the multicast. If the host fails to respond with a Membership Report, the Mrouter stops sending the multicast to that path.
- The host can send a Leave report to the switch, which sends a proxy Leave report to the Mrouter. The multicast path is terminated immediately.

FastLeave

When the switch with IGMP Snooping enabled receives an IGMPv2 leave message, it sends a Group-Specific Query to determine if any other devices in the same group (and on the same port) are still interested in the specified multicast group traffic. The switch removes the affiliated port from that particular group, if the following conditions apply:

- If the switch does not receive an IGMP Membership Report message within the query-response-interval
- If no multicast routers have been learned on that port.

With Fastleave enabled on the VLAN, a port can be removed immediately from the port list of the group entry when the IGMP Leave message is received, unless a multicast router was learned on the port.

Enable FastLeave only on VLANs that have only one host connected to each physical port.

IGMP Filtering

With IGMP Filtering, you can allow or deny a port to send and receive multicast traffic to certain multicast groups. Unauthorized users are restricted from streaming multicast traffic across the network.

If access to a multicast group is denied, IGMP Membership Reports from the port for that group are dropped, and the port is not allowed to receive IP multicast traffic from that group. If access to the multicast group is allowed, Membership Reports from the port are forwarded for normal processing.

To configure IGMP Filtering, you must globally enable IGMP Filtering, define an IGMP Filter, assign the filter to a port, and enable IGMP Filtering on the port. To define an IGMP Filter, you must configure a range of IP multicast groups, choose whether the filter will allow or deny multicast traffic for groups within the range, and enable the filter.

NOTE: Low-numbered filters take precedence over high-number filters. For example, the action defined for IGMP Filter 1 supersedes the action defined for IGMP Filter 2.

Configuring the range

Each IGMP Filter allows you to set a start and end point that defines the range of IP addresses upon which the filter takes action. Each IP address in the range must be between 224.0.0.0 and 239.255.255.255.

Configuring the action

Each IGMP Filter can allow or deny IP multicasts to the range of IP addresses configured. If you configure the filter to deny IP multicasts, then IGMP Membership Reports from multicast groups within the range are dropped.

You can configure a secondary filter to allow IP multicasts to a small range of addresses within a larger range that a primary filter is configured to deny. The two filters work together to allow IP multicasts to a small subset of addresses within the larger range of addresses. The secondary filter must have a lower number than the primary filter, so that it takes precedence.

Static multicast router

A static multicast router (Mrouter) can be configured for a particular port on a particular VLAN. A static Mrouter does not have to be learned through IGMP Snooping.

A total of eight static Mrouters can be configured on the switch. A port that belongs to a trunk group cannot accept a static Mrouter, only Mrouters learned through IGMP Snooping.

When you configure a static Mrouter on a VLAN, it replaces any dynamic Mrouters learned through IGMP Snooping.

IGMP Snooping configuration example

This section provides steps to configure IGMP Snooping on the switch, using the Command Line Interface (CLI) or the Browser-based Interface (BBI).

Configuring IGMP Snooping (CLI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (CLI example)” section in the “VLANs” chapter.
2. Add VLANs to IGMP Snooping and enable the feature.

```
>> /cfg/l3/igmp/snoop                (Select IGMP Snooping menu)
>> IGMP Snoop# ena                   (Enable IGMP Snooping)
>> IGMP Snoop# apply                 (Make your changes active)
```

3. View dynamic IGMP information.

```
>> /info/l3/igmp                  (Select IGMP Information menu)
>> IGMP Multicast# dump           (Show IGMP Group information)

>> Switch-A - IGMP Multicast# dump
      Group      VLAN      Version          Port
  -----------  ------    ---------     -------------
   238.1.0.0        1       V2              20
   238.1.0.1        1       V2              21
>> IGMP Multicast# mrouter        (Select MRouter Information menu)
>> IGMP Multicast Router# dump    (Show IGMP Group information)

      VLAN           Port              Version     L)earnt/(S)tatic
  ----------     -----------      ----------   -----------------
       1              23                 V2             S
```

These commands display information about IGMP Groups and Mrouters learned through IGMP Snooping.

Configuring IGMP Filtering (CLI example)

1. Enable IGMP Filtering on the switch.

```
        >> /cfg/l3/igmp/igmpflt             (Select IGMP Filtering menu)
        >> IGMP Filter# ena                  (Enable IGMP Filtering)
        Current status: disabled
        New status: enabled
```

2. Define an IGMP Filter.

```
>> //cfg/l3/igmp/igmpflt               (Select IGMP Filtering menu)
>>IGMP Filter# filter 1                (Select Filter 1 Definition menu)
>>IGMP Filter 1 Definition# range 224.0.1.0  (Enter first IP
                                              address of the range)
Current multicast address2:
Enter new multicast address2: 226.0.0.0      (Enter second IP
                                              address of the range)
Current multicast address1:
New pending multicast address1: 224.0.1.0
Current multicast address2:
New pending multicast address2: 226.0.0.0
>>IGMP Filter 1 Definition# action deny  (Deny multicast traffic)
>>IGMP Filter 1 Definition# ena          (Enable the filter)
```

3. Assign the IGMP Filter to a port.

```
>> //cfg/l3/igmp/igmpflt          (Select IGMP Filtering menu)
>>IGMP Filter# port 24             (Select port 24)
>>IGMP Port 24# filt ena      (Enable IGMP Filtering on the port)
Current port 24 filtering: disabled
New port 24 filtering: enabled
>>IGMP Port 24# add 1              (Add IGMP Filter 1 to the port)
>>IGMP Port 24# apply              (Make your changes active
```

Configuring a Static Mrouter (CLI example)

1. Configure a port to which the static Mrouter is connected, and enter the appropriate VLAN.

```
>> /cfg/l3/igmp/mrouter                (Select IGMP Mrouter menu)
>> Static Multicast Router# add 20   (Add port 20 as Static
                                       Mrouter port)
Enter VLAN number: (1-4094) 1         (Enter the VLAN number)
Enter the version number of mrouter [1|2]: 2 (Enter the IGMP
                                              version number)
```

2. Apply, verify, and save the configuration.

```
>> Static Multicast Router# apply     (Apply the configuration)
>> Static Multicast Router# cur       (View the configuration)
>> Static Multicast Router# save      (Save the configuration)
```

Configuring IGMP Snooping (BBI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (BBI example)” section in the “VLANs” chapter.
2. Configure IGMP Snooping.
   a. Click the Configure context button.
   b. Open the IGMP folder, and select IGMP Snooping (click the underlined text, not the folder).

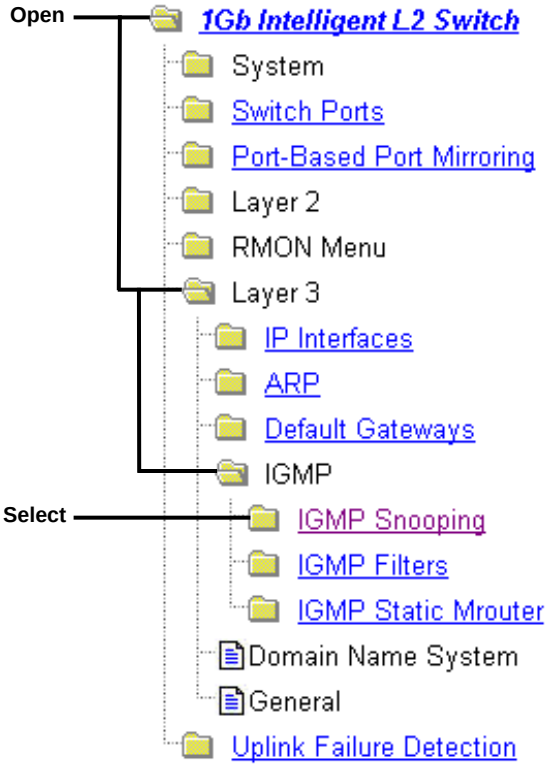

c. Enable IGMP Snooping.

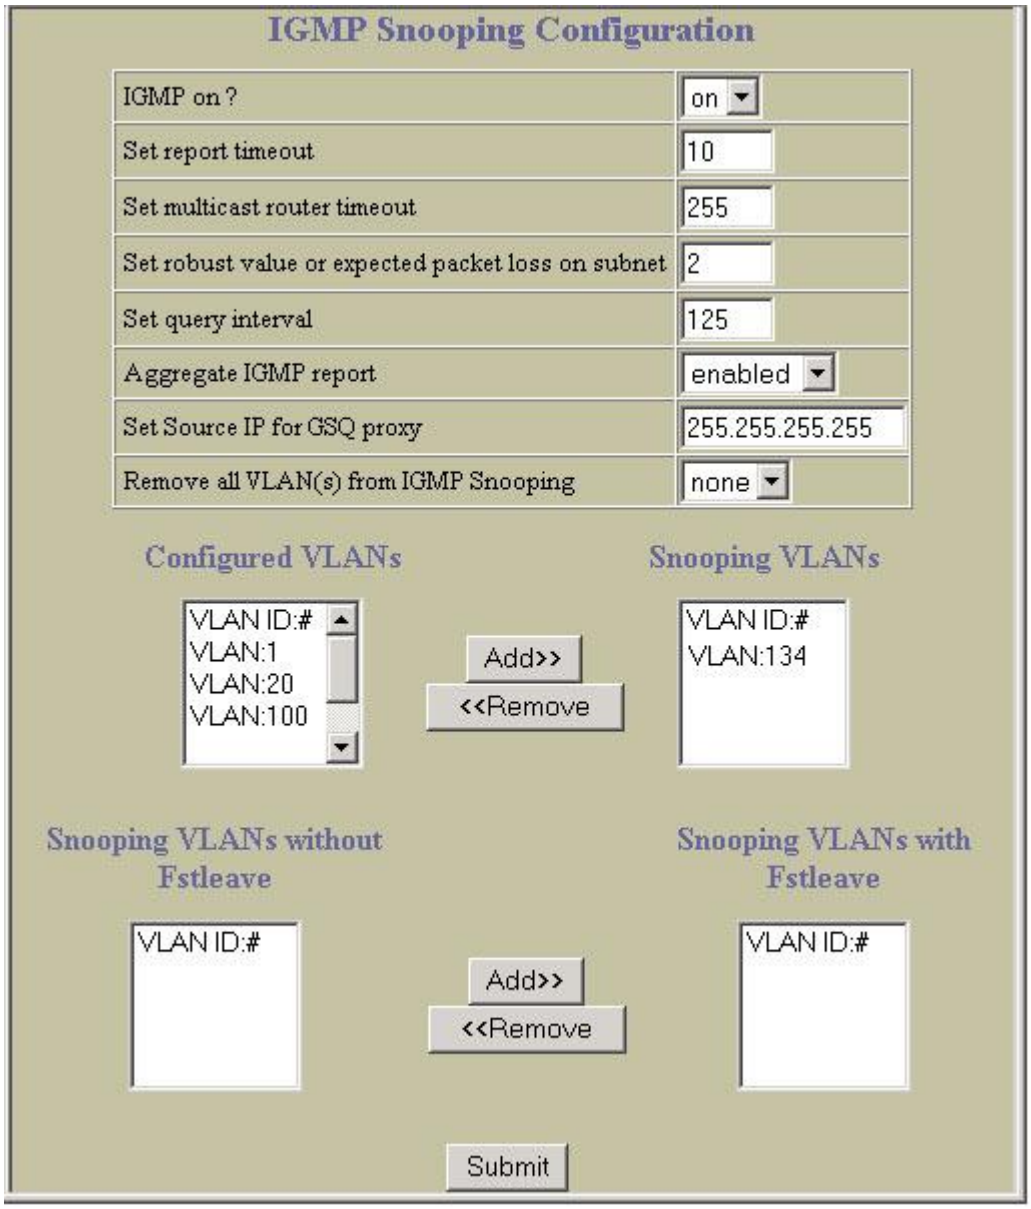

d. Click Submit.

3. Apply, verify, and save the configuration.

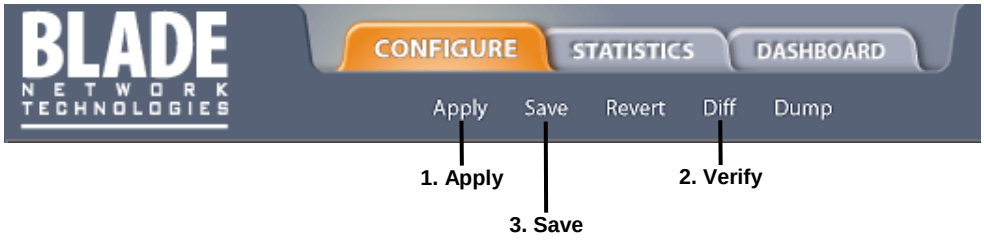

Configuring IGMP Filtering (BBI example)

1. Configure IGMP Snooping.
2. Enable IGMP Filtering.
   a. Click the Configure context button.
   b. Open the IGMP folder, and select IGMP Filters (click the underlined text, not the folder).

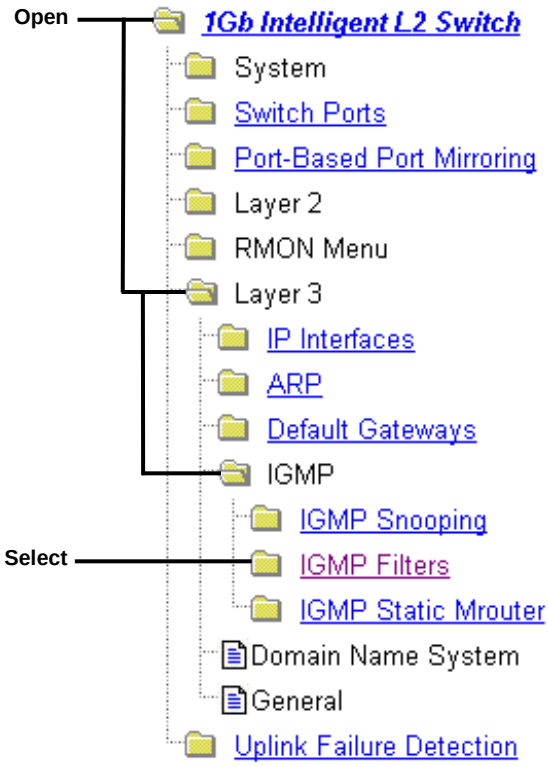

   c. Enable IGMP Filtering globally.

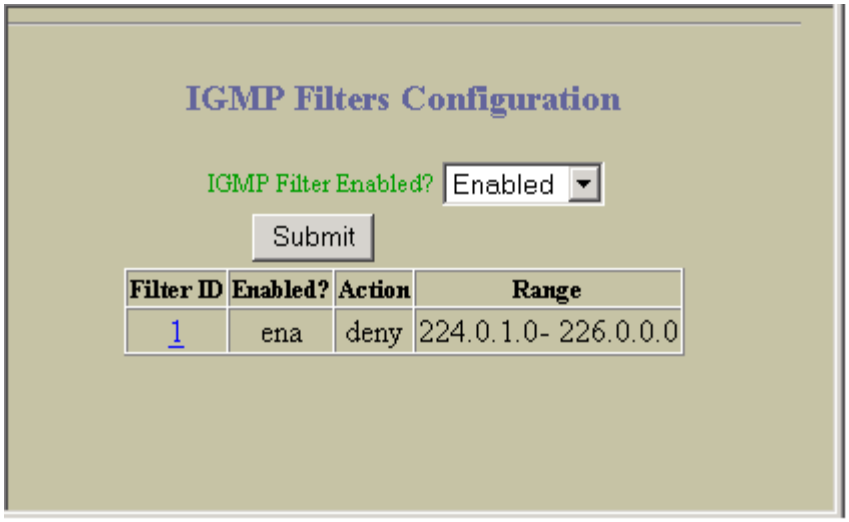

   d. Click Submit.

3. Define the IGMP Filter.

   a. Select Layer 3 > IGMP > IGMP Filters > Add Filter.

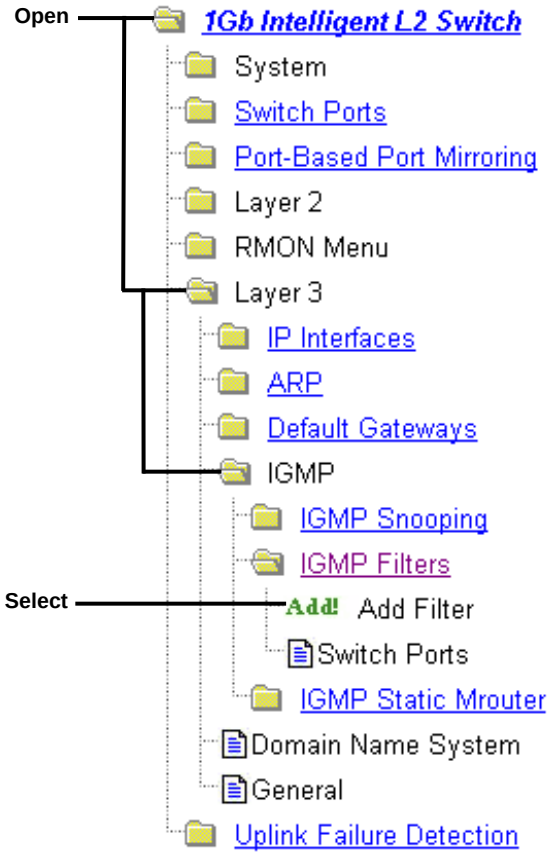

   b. Enable the IGMP Filter. Assign the range of IP multicast addresses and the filter action (allow or deny).

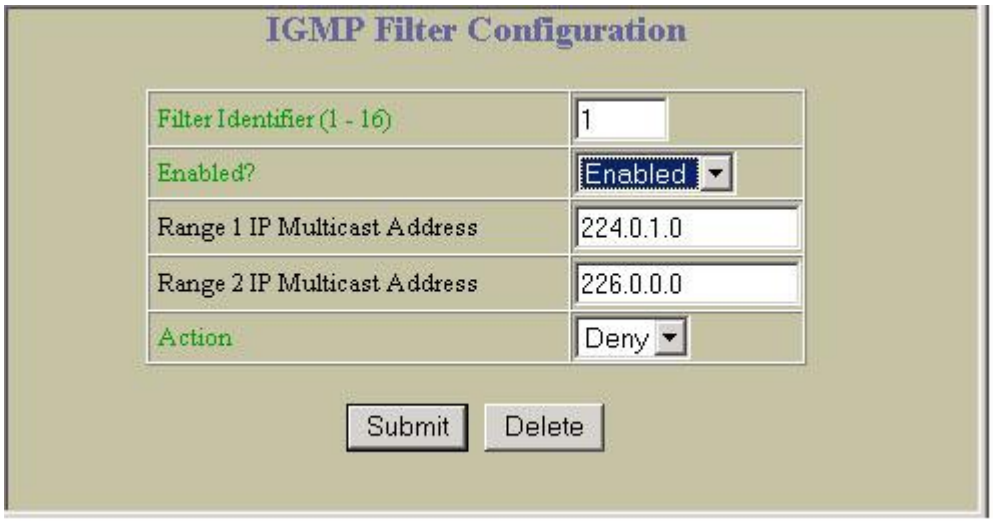

   c. Click Submit.

4. Assign the filter to a port and enable IGMP Filtering on the port.
   a. Select Layer 3 > IGMP > IGMP Filters > Switch Ports.

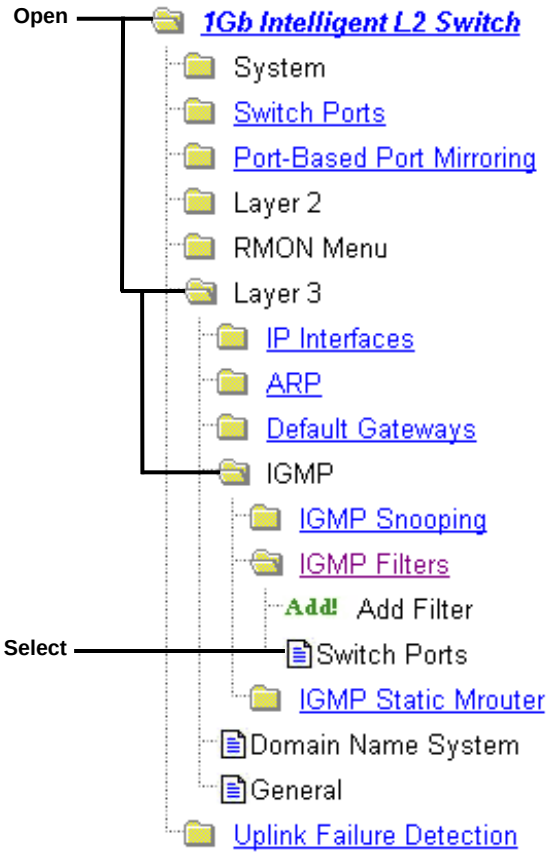

   b. Select a port from the list.

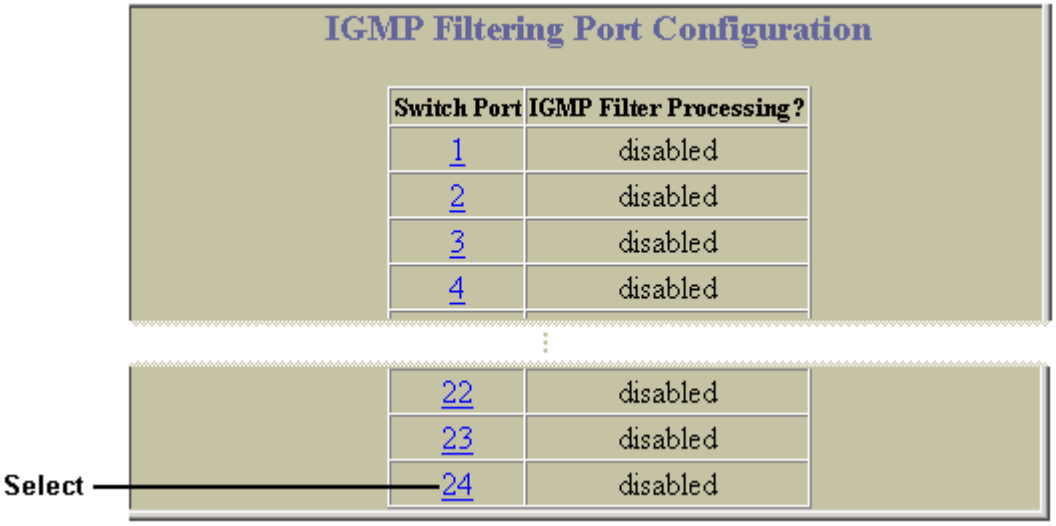

c. Enable IGMP Filtering on the port. Select a filter in the IGMP Filters Available list, and click Add.

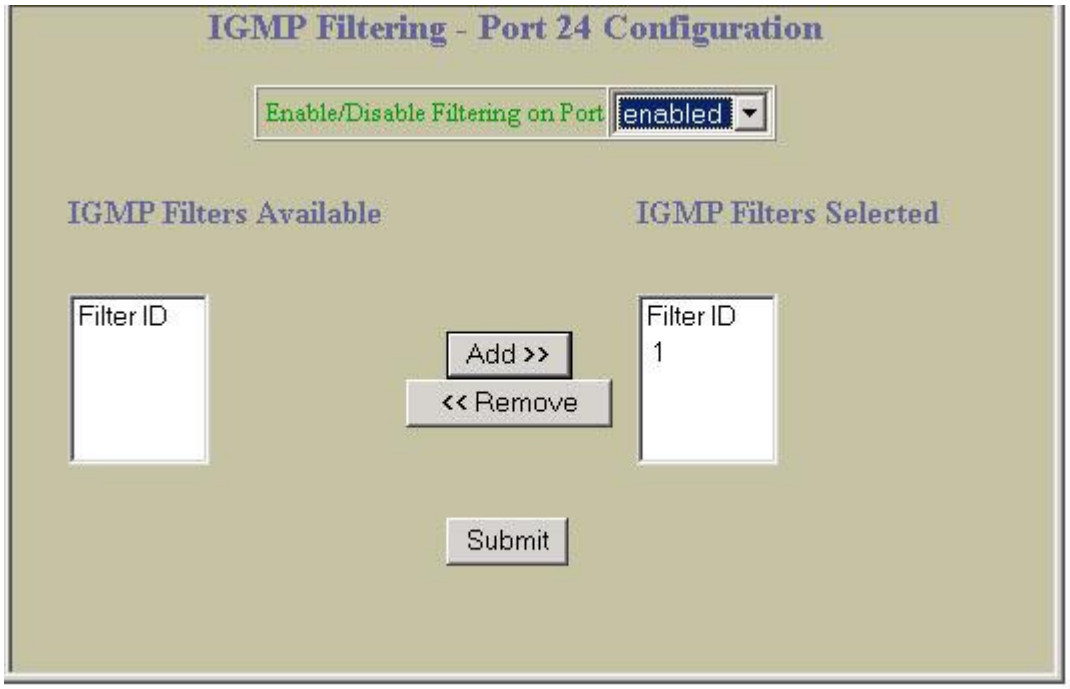

d. Click Submit.

5. Apply, verify, and save the configuration.

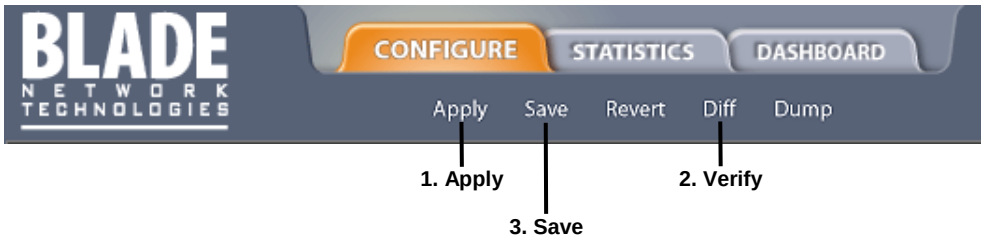

Configuring a Static Multicast Router (BBI example)

1. Configure Static Mrouter.
   a. Click the Configure context button.
   b. Open the Switch folder and select IP Menu > IGMP > IGMP Static MRouter.
   c. Enter a port number, VLAN ID number, and IGMP version number.

   d. Click Submit.

2. Apply, verify, and save the configuration.

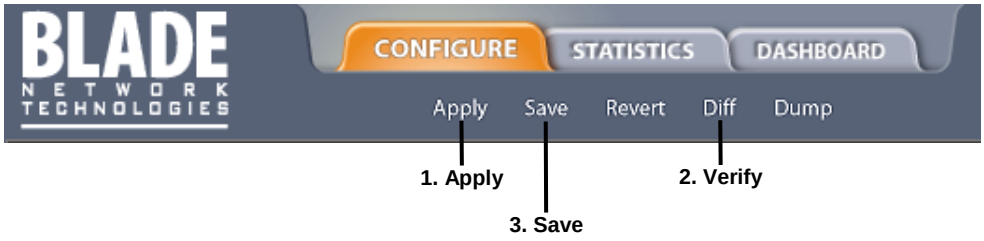

Remote monitoring

Introduction

Remote Monitoring (RMON) allows network devices to exchange network monitoring data.

RMON performs the following major functions:

- Gathers cumulative statistics for Ethernet interfaces
- Tracks a history of statistics for Ethernet interfaces
- Creates and triggers alarms for user-defined events

Overview

The RMON MIB provides an interface between the RMON agent on the switch and an RMON management application. The RMON MIB is described in RFC 1757.

The RMON standard defines objects that are suitable for the management of Ethernet networks. The RMON agent continuously collects statistics and proactively monitors switch performance. RMON allows you to monitor traffic flowing through the switch.

The switch supports the following RMON Groups, as described in RFC 1757:

- Group 1: Statistics
- Group 2: History
- Group 3: Alarms
- Group 9: Events

RMON group 1 — statistics

The switch supports collection of Ethernet statistics as outlined in the RMON statistics MIB, in reference to etherStatsTable. You can enable RMON statistics on a per-port basis, and you can view them using the `/stat/port x/rmon` command. RMON statistics are sampled every second, and new data overwrites any old data on a given port.

NOTE: RMON port statistics must be enabled for the port before you can view RMON statistics.

Configuring RMON Statistics (AOS CLI example)

1. Enable RMON on each port where you wish to collect RMON statistics.

```
>> /cfg/port 23/rmon                    (Select Port 23 RMON)
>> Port 23 RMON# ena                    (Enable RMON)
>> Port 23 RMON# apply                  (Make your changes active)
>> Port 23 RMON# save                   (Save for restore after reboot)
```

2. View RMON statistics for the port.

```
>> /stats/port 23                                (Select Port 23 Stats)
>> Port Statistics# rmon
------------------------------------------------------------------------

RMON statistics for port 23:
etherStatsDropEvents:                          NA
etherStatsOctets:                         7305626
etherStatsPkts:                             48686
etherStatsBroadcastPkts:                     4380
etherStatsMulticastPkts:                     6612
etherStatsCRCAlignErrors:                      22
etherStatsUndersizePkts:                        0
etherStatsOversizePkts:                         0
etherStatsFragments:                            2
etherStatsJabbers:                              0
etherStatsCollisions:                           0
etherStatsPkts64Octets:                     27445
etherStatsPkts65to127Octets:                12253
etherStatsPkts128to255Octets:                1046
etherStatsPkts256to511Octets:                 619
etherStatsPkts512to1023Octets:               7283
etherStatsPkts1024to1518Octets:                38
```

Configuring RMON Statistics (BBI example)

1. Configure ports.
   a. Click the Configure context button.
   b. Select Switch Ports (click the underlined text, not the folder).

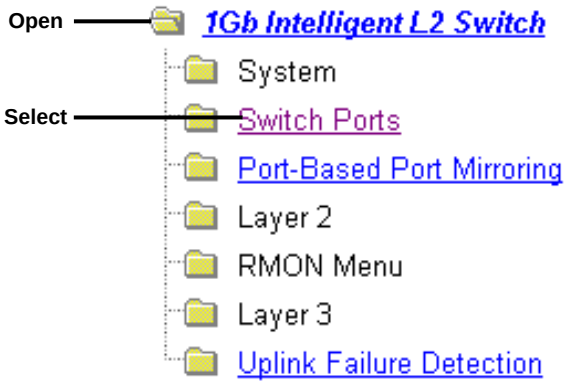

2. Select a port.

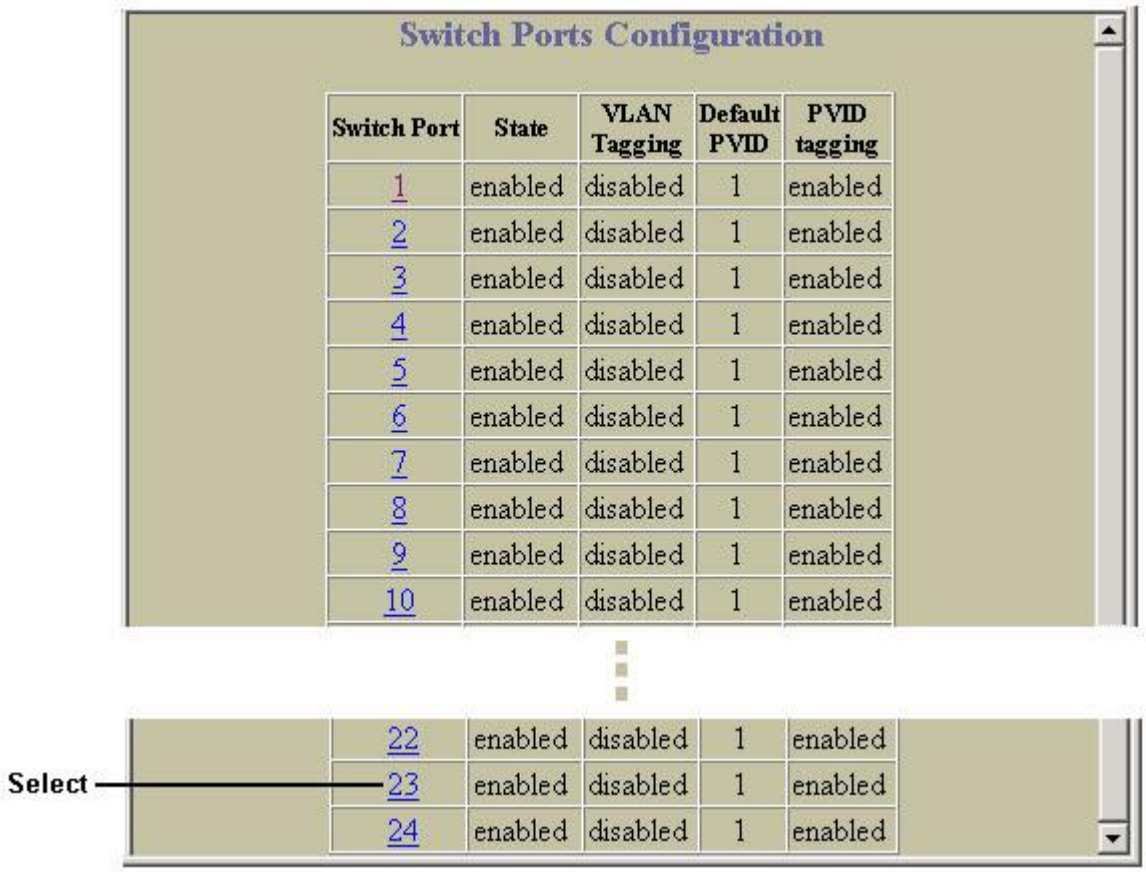

3. Enable RMON on the port.

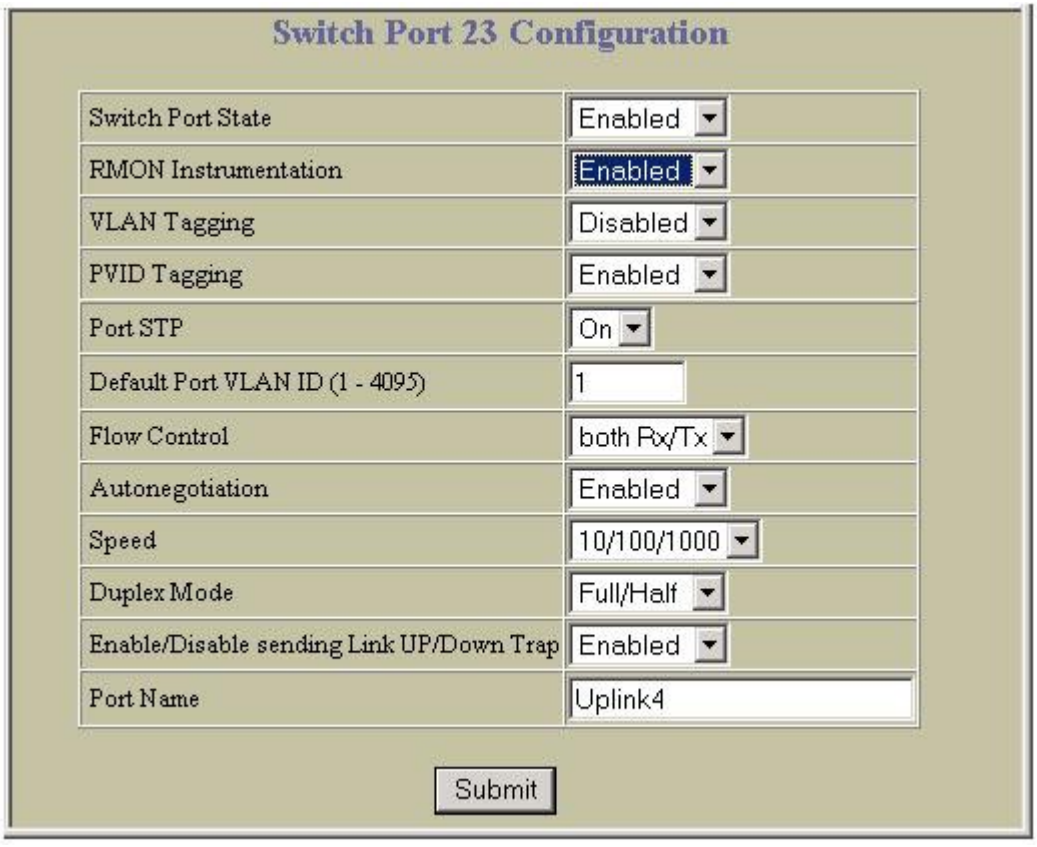

4. Click Submit.
5. Apply, verify, and save the configuration.

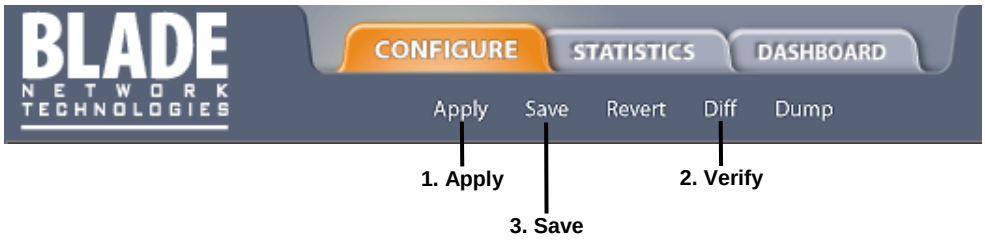

RMON group 2 — history

The RMON History group allows you to sample and archive Ethernet statistics for a specific interface during a specific time interval. The switch supports up to five RMON History groups.

NOTE: RMON port statistics must be enabled for the port before an RMON history group can monitor the port.

Data is stored in buckets, which store data gathered during discreet sampling intervals. At each configured interval, the history instance takes a sample of the current Ethernet statistics, and places them into a bucket. History data buckets reside in dynamic memory. When the switch is re-booted, the buckets are emptied.

Requested buckets (`/cfg/rmon/hist x/rbnum`) are the number of buckets, or data slots, requested by the user for each History Group. Granted buckets (`/info/rmon/hist x/gbnum`) are the number of buckets granted by the system, based on the amount of system memory available. The system grants a maximum of 50 buckets.

Use an SNMP browser to view History samples.

History MIB objects

The type of data that can be sampled must be of an ifIndex object type, as described in RFC1213 and RFC1573. The most common data type for the history sample is as follows:

`1.3.6.1.2.1.2.2.1.1.x -mgmt.interfaces.ifTable.ifIndex.interface`

The last digit (*x*) represents the interface on which to monitor, which corresponds to the port number (1-24). History sampling is done per port, by utilizing the interface number to specify the port number.

Configure RMON History (AOS CLI example)

1. Enable RMON on each port where you wish to collect RMON History.

```
>> /cfg/port 23/rmon                    (Select Port 23 RMON)
>> Port 23# ena                         (Enable RMON)
>> Port 23 RMON# apply                  (Make your changes active)
>> Port 23 RMON# save                   (Save for restore after reboot)
```

2. Configure the RMON History parameters.

```
>> /cfg/rmon/hist 1                          (Select RMON History 1)
>> RMON History 1# ifoid 1.3.6.1.2.1.2.2.1.1.23
>> RMON History 1# rbnum 30
>> RMON History 1# intrval 120
>> RMON History 1# owner "Owner_History_1"
```

3. Apply and save the configuration.

```
>> RMON History 1# apply              (Make your changes active)
>> RMON History 1# save               (Save for restore after reboot)
```

This configuration creates an RMON History group to monitor port 23. It takes a data sample every two minutes, and places the data into one of the 30 requested buckets. After 30 samples are gathered, the new samples overwrite the previous samples, beginning with the first bucket. Use SNMP to view the data.

Configure RMON History (BBI example)

1. Configure an RMON History group.
   a. Click the Configure context button.
   b. Open the Switch folder, and select RMON > History > Add History Group.

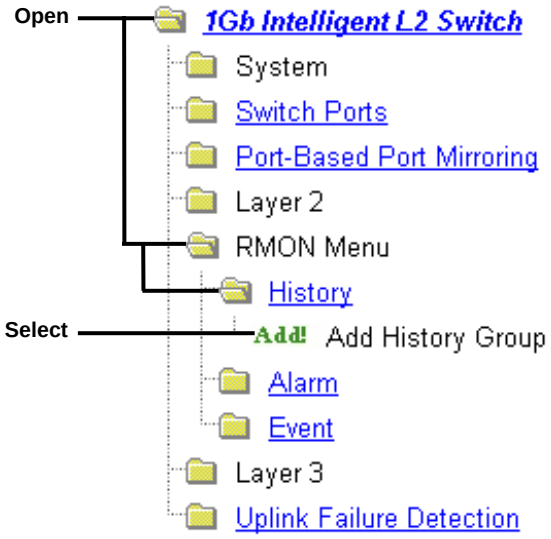

2. Configure RMON History Group parameters.

3. Click Submit.
4. Apply, verify, and save the configuration.

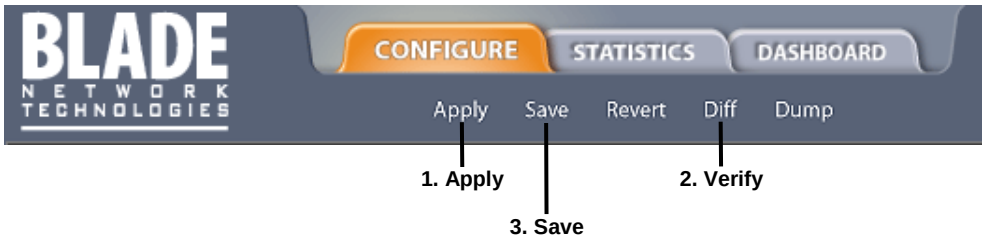

RMON group 3 — alarms

The RMON Alarm group allows you to define a set of thresholds used to determine network performance. When a configured threshold is crossed, an alarm is generated. For example, you can configure the switch to issue an alarm if more than 1,000 CRC errors occur during a 10-minute time interval. The switch supports up to 30 RMON Alarm groups.

Each Alarm index consists of a variable to monitor, a sampling time interval, and parameters for rising and falling thresholds. The Alarm group can be used to track rising or falling values for a MIB object. The object must be a counter, gauge, integer, or time interval.

Use the `/cfg/rmon/alarm x/revtidx` or `/fevtidx` to correlate an alarm index to an event index. When the alarm threshold is reached, the corresponding event is triggered.

Alarm MIB objects

The most common data types used for alarm monitoring are ifStats: errors, drops, bad CRCs, and so on. These MIB Object Identifiers (OIDs) correlate to the ones tracked by the History group. An example of an ICMP stat is as follows:

`1.3.6.1.2.1.5.1.0 – mgmt.icmp.icmpInMsgs`

The last digit (*x*) represents the interface on which to monitor, which corresponds to the interface number, or port number, as follows:

1-256 = IF 1-256
257 = port 1
258 = port 2
…
280 = port 24

This value represents the alarm's MIB OID, as a string. Note that for non-tables, you must supply a .0 to specify end node.

Configure RMON Alarms (AOS CLI example 1)

1. Configure the RMON Alarm parameters to track the number of packets received on a port.

```
>> /cfg/rmon/alarm 6                                (Select RMON Alarm 6)
>> RMON Alarm 6# oid 1.3.6.1.2.1.2.2.1.10.276
>> RMON Alarm 6# intrval 3600
>> RMON Alarm 6# almtype rising
>> RMON Alarm 6# rlimit 2000000000
>> RMON Alarm 6# revtidx 6
>> RMON Alarm 6# sample abs
>> RMON Alarm 6# owner "Alarm_for_ifInOctets"
```

2. Apply and save the configuration.

```
>> RMON Alarm 6# apply                (Make your changes active)
>> RMON Alarm 6# save                 (Save for restore after reboot)
```

This configuration creates an RMON alarm that checks `ifInOctets` on port 20 once every hour. If the statistic exceeds two billion, an alarm is generated that triggers event index 6.

Configure RMON Alarms (AOS CLI example 2)

1. Configure the RMON Alarm parameters to track ICMP messages.

```
>> /cfg/rmon/alarm 5                              (Select RMON Alarm 5)
>> RMON Alarm 5# oid 1.3.6.1.2.1.5.8.0
>> RMON Alarm 5# intrval 60
>> RMON Alarm 5# almtype rising
>> RMON Alarm 5# rlimit 200
>> RMON Alarm 5# revtidx 5
>> RMON Alarm 5# sample delta
>> RMON Alarm 5# owner "Alarm_for_icmpInEchos"
```

2. Apply and save the configuration.

```
>> RMON Alarm 5# apply                 (Make your changes active)
>> RMON Alarm 5# save                  (Save for restore after reboot)
```

This configuration creates an RMON alarm that checks `icmpInEchos` on the switch once every minute. If the statistic exceeds 200 within a 60 second interval, an alarm is generated that triggers event index 5.

Configure RMON Alarms (BBI example 1)

1. Configure an RMON Alarm group.
   a. Click the Configure context button.
   b. Open the Switch folder, and select RMON > Alarm > Add Alarm Group.

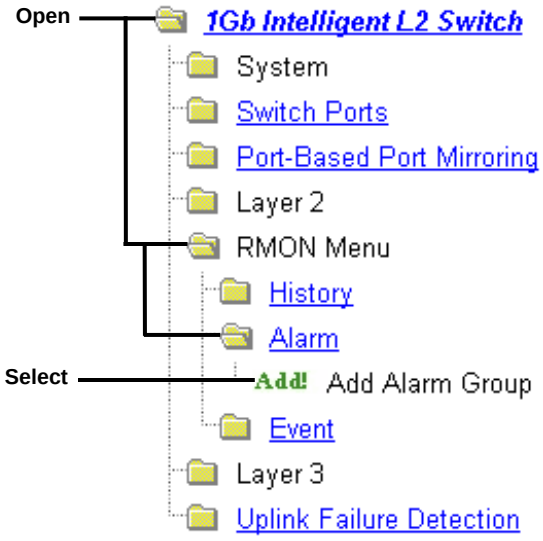

Configure RMON Alarm Group parameters to check `ifInOctets` on port 19 once every hour. Enter a rising limit of two billion, and a rising event index of 6. This configuration creates an RMON alarm that checks `ifInOctets` on port 19 once every hour. If the statistic exceeds two billion, an alarm is generated that triggers event index 6.

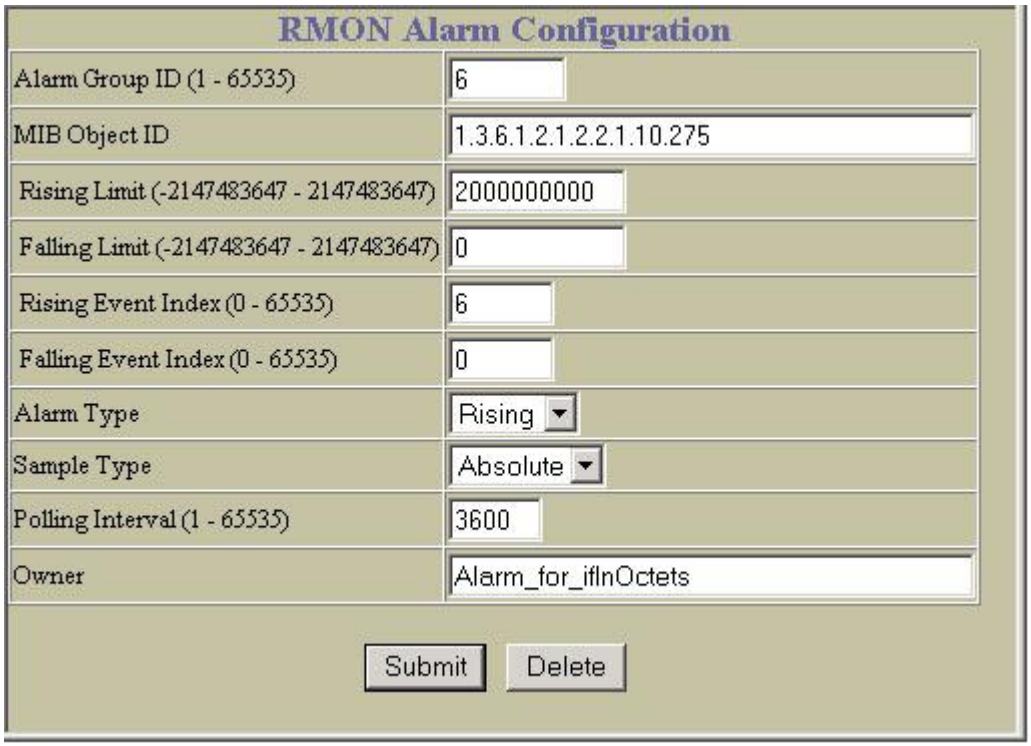

2. Click Submit.
3. Apply, verify, and save the configuration.

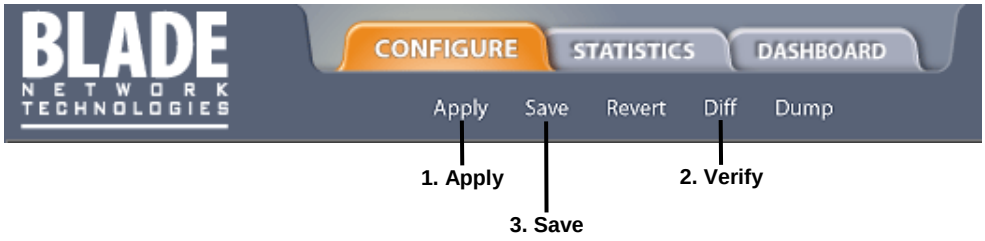

Configure RMON Alarms (BBI example 2)

1. Configure an RMON Alarm group.
   a. Click the Configure context button.
   b. Open the Switch folder, and select RMON > Alarm > Add Alarm Group.

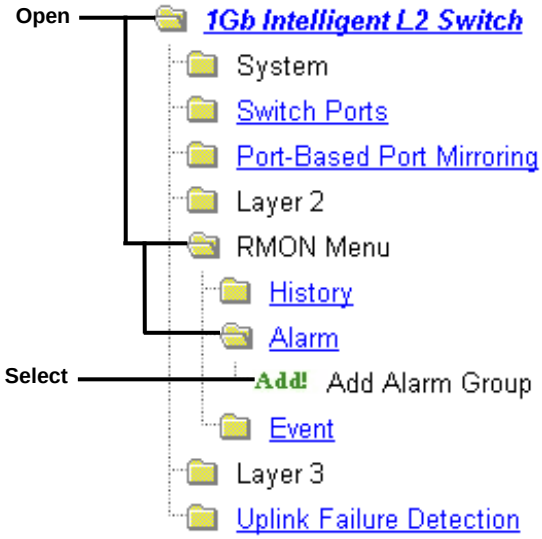

Configure RMON Alarm Group parameters to check `icmpInEchos`, with a polling interval of 60, a rising limit of 200, and a rising event index of 5. This configuration creates an RMON alarm that checks `icmpInEchos` on the switch once every minute. If the statistic exceeds 200 within a 60 second interval, an alarm is generated that triggers event index 5.

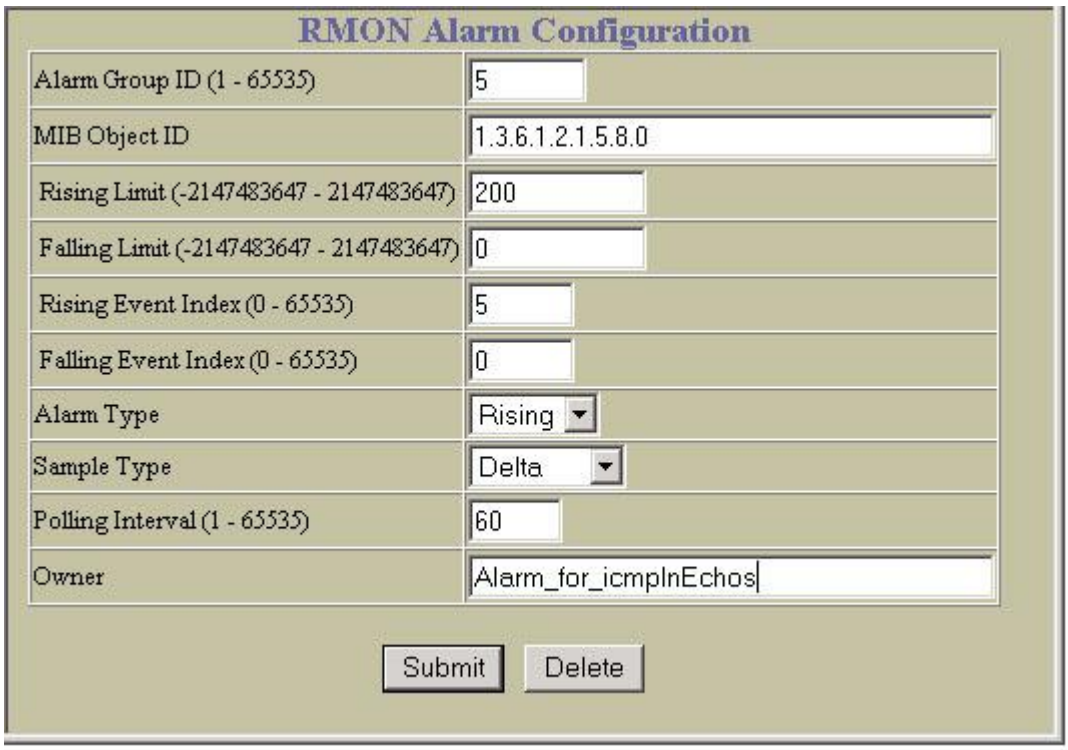

2. Click Submit.

3. Apply, verify, and save the configuration.

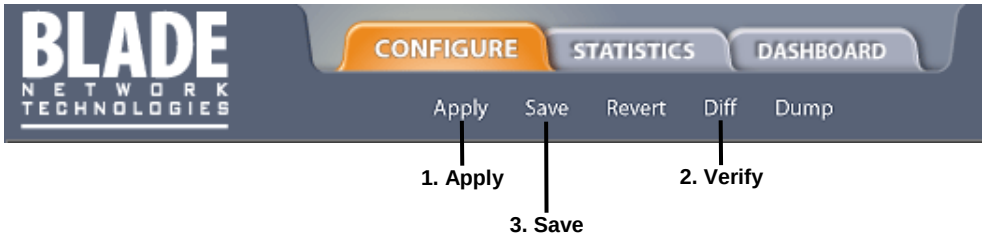

RMON group 9 — events

The RMON Event group allows you to define events that are triggered by alarms. An event can be a log message, an SNMP trap message, or both. The switch supports up to 30 RMON Event groups.

When an alarm is generated, it triggers a corresponding event notification. Use the `/cfg/rmon/alarm x/revtidx` and `/fevtidx` commands to correlate an event index to an alarm.

RMON events use SNMP and syslogs to send notifications. Therefore, an SNMP trap host must be configured for trap event notification to work properly.

RMON uses a SYSLOG host to send syslog messages. Therefore, an existing SYSLOG host (`/cfg/sys/syslog`) must be configured for event log notification to work properly. Each log event generates a SYSLOG of type RMON that corresponds to the event.

Configuring RMON Events (AOS CLI example)

1. Configure the RMON Event parameters.

```
>> /cfg/rmon/event 5                              (Select RMON Event 5)
>> RMON Event 5# descn "SYSLOG_generation_event"
>> RMON Event 5# type log
>> RMON Event 5# owner "Owner_event_5"
```

2. Apply and save the configuration.

```
>> RMON Alarm 5# apply                 (Make your changes active)
>> RMON Alarm 5# save                  (Save for restore after reboot)
```

This configuration creates an RMON event that sends a SYSLOG message each time it is triggered by an alarm.

Configuring RMON Events (BBI example)

1. Configure an RMON Event group.
   a. Click the Configure context button.
   b. Open the Switch folder, and select RMON > Event > Add Event Group.

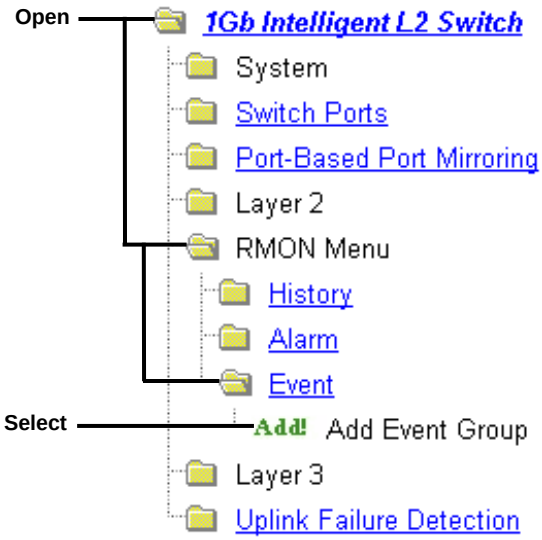

Configure RMON Event Group parameters. This configuration creates an RMON event that sends a SYSLOG message each time it is triggered by an alarm.

2. Click Submit.
3. Apply, verify, and save the configuration.

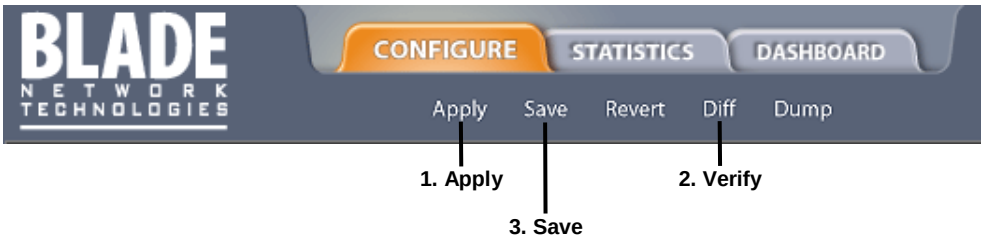

High availability

Introduction

Switches support high availability network topologies. This release provides information about Uplink Failure Detection.

Uplink Failure Detection

Uplink Failure Detection (UFD) is designed to support Network Adapter Teaming on the CPU Blades.

UFD allows the switch to monitor specific uplink ports to detect link failures. When the switch detects a link failure, it automatically disables specific downlink ports. The corresponding server's network adapter can detect the disabled downlink, and trigger a network-adapter failover to another port on the switch, or another switch in the chassis.

The switch automatically enables the downlink ports when the uplink returns to service.

The following figure shows a basic UFD configuration, with a Failure Detection Pair (FDP) that consists of one LtM (Link to Monitor) and one LtD (Link to Disable). When the switch detects a link failure in the LtM, it disables the ports in the LtD. The server blade detects the disabled downlink port, which triggers a NIC failover.

Figure 10 Uplink Failure Detection for switches

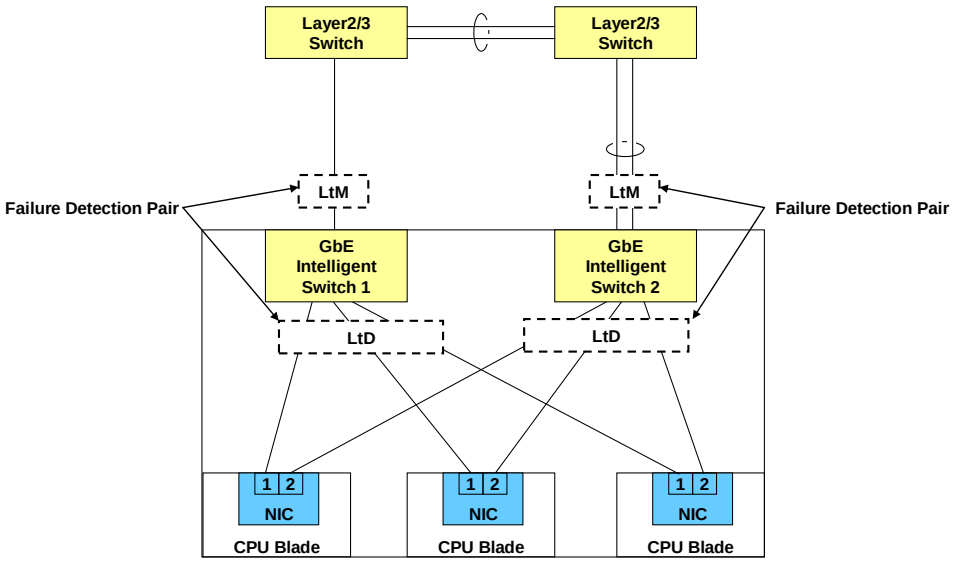

NOTE: The port numbers specifie Blade Enclosure ıay not directly correspond to the physical port configuration of your system.

Failure Detection Pair

To use UFD, you must configure a Failure Detection Pair and then turn UFD on. A Failure Detection Pair consists of the following groups of ports:

- Link to Monitor (LtM)
  The Link to Monitor group consists of one uplink port (20-24), one trunk group that contains only uplink ports, or one LACP trunk group that contains only uplink ports. The switch monitors the LtM for link failure.
- Link to Disable (LtD)
  The Link to Disable group consists of one or more downlink ports (1-16) and trunk groups that contain only downlink ports and LACP trunk groups that contain only downlink ports. When the switch detects a link failure on the LtM, it automatically disables all ports in the LtD.
- When the LtM returns to service, the switch automatically enables all ports in the LtD.

Spanning Tree Protocol with UFD

If Spanning Tree Protocol (STP) is enabled on ports in the LtM, then the switch monitors the STP state and the link status on ports in the LtM. The switch automatically disables the ports in the LtD when it detects a link failure or STP Blocking state.

When the switch determines that ports in the LtM are in STP Forwarding State, then it automatically enables the ports in the LtD, to fall back to normal operation.

Configuration guidelines

This section provides important information about configuring UFD:

- UFD is required only when uplink-path redundancy is not available on the blade switches.
- Four Failure Detection pairs (one group of Links to Monitor and one group of Links to Disable) is supported on each switch (all VLANs and Spanning Tree Groups).
- An LtM can be any one of one uplink port, one trunk group of uplink ports, or one LACP trunk group of uplink ports. Ports that are already members of a trunk group are not allowed to be assigned to an LtM.
- A trunk group or a LACP trunk group configured as an LtM can contain multiple uplink ports (20-24), but no downlink ports (1-16) or interconnect ports (17-18). An uplink port cannot be added to a trunk group if it already belongs to an LtM.
- An LtD can contain one or more ports, and/or one or more trunks, and/or one or more LACP trunks.
- A trunk group or a LACP trunk group configured as an LtD can contain multiple downlink ports (1-16), but no uplink ports (20-24) or interconnect ports (17-18).

Monitoring Uplink Failure Detection

The UFD information menu displays the current status of the LtM and LtD, and their member ports or trunks. For example:

```
>> Information# ufd
Uplink Failure Detection 1: Enabled
LtM status: Down
Member      STG    STG State       Link Status
---------   ---    ------------    -----------
  port 24                           down
             1       DISABLED
            10       DISABLED *
            15       DISABLED *
* = STP turned off for this port.

LtD status: Auto Disabled
Member       Link Status
---------    -----------
  port  1      disabled
  port  2      disabled
  port  3      disabled
  port  4      disabled

Uplink Failure Detection 2: Disabled

Uplink Failure Detection 3: Disabled

Uplink Failure Detection 4: Disabled
```

Use the `/stats/ufd` command to find out how many times link failure was detected on the LtM, how many times Spanning Tree blocking state was detected on the LtM, and how many times UFD disabled ports in the LtD.

Configuring Uplink Failure Detection

The preceding figure shows a basic UFD configuration. Port 21 on Blade Switch 1 is connected to a Layer 2/3 routing switch outside of the chassis. Port 20 and port 22 on Blade Switch 2 form a trunk that is connected to a different Layer 2/3 routing switch. The interconnect ports (17-18) are disabled.

In this example, the port 1 of a NIC is the primary network adapter, the port 2 of the NIC is a non-primary adapter. The port 1 of the NIC on the CPU blade server 1 and the CPU blade server 2 are connected to port 1 and port 2 on the Switch 1. The port 2 of the NIC on the CPU blade server 1 and the CPU blade server 2 are connected to port 1 and port 2 on the Switch 2.

Configuring UFD on Switch 1 (AOS CLI example)

1. Assign uplink ports (20-24) to be monitored for communication failure.

```
>> Main# /cfg/ufd/fdp 1/ena           (Enable Failure Detection Pair 1)
>> FDP# ltm                           (Select Link to Monitor menu)
>> Failure Link to Monitor# addport 21 (Monitor uplink port 21)
```

2. Assign downlink ports (1-16) to disable when an uplink failure occurs.

```
>> /cfg/ufd/fdp 1/ltd                  (Select Link to Disable menu)
>> Failure Link to Disable# addport 1  (Add port 1 as a Link to Disable)
>> Failure Link to Disable# addport 2  (Add port 2 as a Link to Disable)
```

3. Turn UFD on.

```
>> /cfg/ufd/on                        (Turn Uplink Failure Detection on)
>> Uplink Failure Detection# apply    (Make your changes active)
>> Uplink Failure Detection# save     (Save for restore after reboot)
```

When a link failure or Spanning Tree blocking occurs on port 21, Switch 1 disables port 1 and port 2.

Configuring UFD on Switch 2 (AOS CLI example)

1. Create a trunk group of uplink ports (20-24) to monitor. First you must set each port to full duplex mode.

```
>> Main# /cfg/port 20/gig/mode full      (Set port 20 to full duplex)
>> Main# /cfg/port 22/gig/mode full      (Set port 22 to full duplex)
>> Main# /cfg/trunk 2                    (Create trunk group 2)
>> Trunk group 2# ena                    (Enable trunk group 2)
>> Trunk group 2# add 20                 (Add port 20 to trunk group 2)
>> Trunk group 2# add 22                 (Add port 22 to trunk group 2)
```

2. Assign the trunk group to be monitored for communication failure.

```
>> Main# /cfg/ufd/fdp 1/ena           (Enable Failure Detection Pair 1)
>> FDP# ltm                              (Select Link to Monitor menu)
>> Failover Link to Monitor# addtrnk 2   (Monitor trunk group 2)
```

3. Assign downlink ports (1-16) to disable when an uplink failure occurs.

```
>> Main# /cfg/ufd/fdp 1/ltd             (Select Link to Disable menu)
>> Failover Link to Disable# addport 1   (Add port 1 as a Link to Disable)
>> Failover Link to Disable# addport 2   (Add port 2 as a Link to Disable)
```

4. Turn UFD on.

```
>> Main# /cfg/ufd/on                 (Turn Uplink Failure Detection on)
>> Uplink Failure Detection# apply (Make your changes active)
>> Uplink Failure Detection# save   (Save for restore after reboot)
```

When a link failure or Spanning Tree blocking occurs on trunk group 2, Switch 2 disables port 1 and port 2.

Configuring Uplink Failure Detection (BBI example)

1. Configure Uplink Failure Detection.
   a. Click the Configure context button.
   b. Open the Switch folder, and select Uplink Failure Detection (click the underlined text, not the folder).

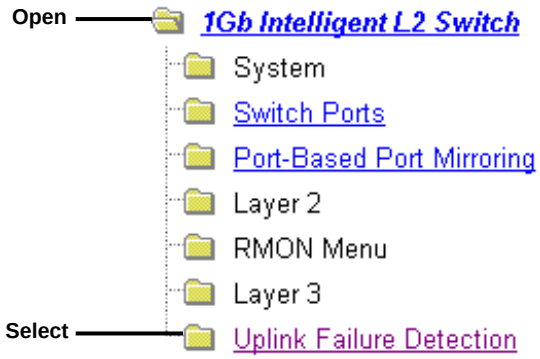

   c. Turn Uplink Failure Detection on, and then select FDP 1.

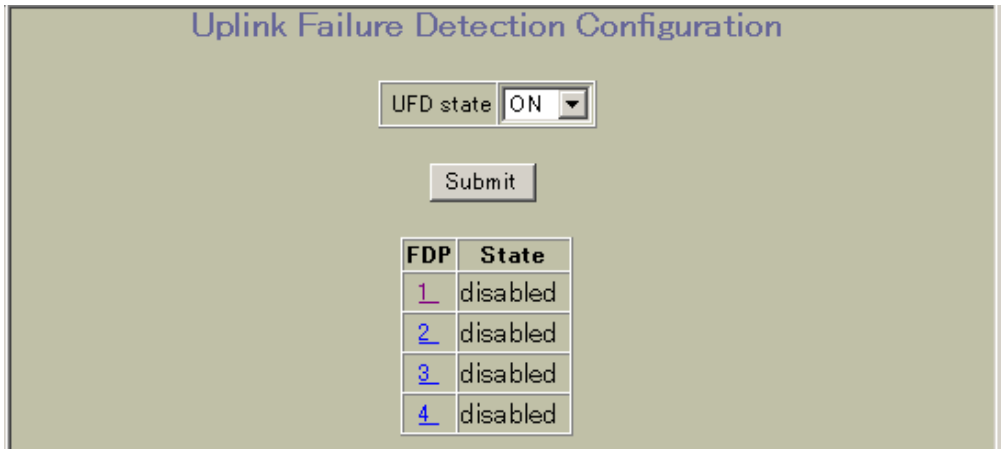

d. Enable the FDP. Select ports in the LtM Ports Available list, and click Add to place the ports into the Link to Monitor (LtM). Select ports in the LtD Ports Available list, and click Add to place the ports into the Link to Disable (LtD).

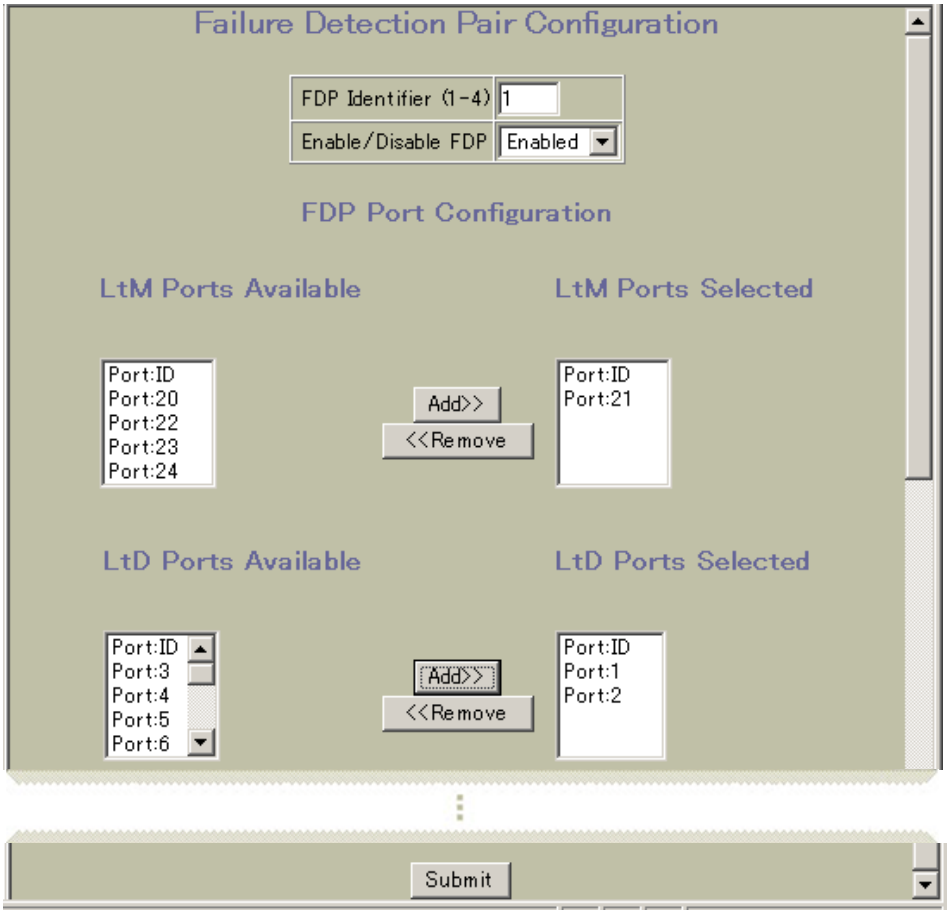

e. Click Submit.

2. Apply, verify, and save the configuration.

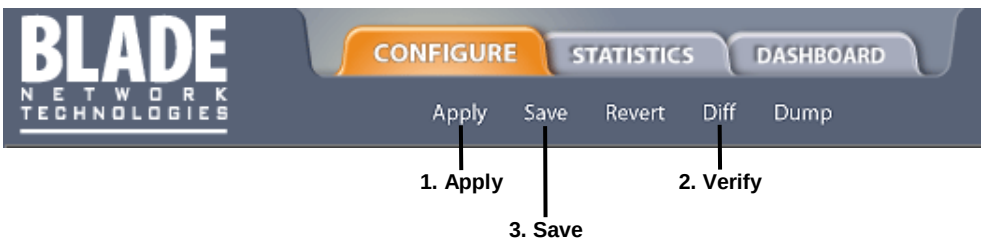

Troubleshooting tools

Introduction

This appendix discusses some tools to help you use the Port Mirroring feature to troubleshoot common network problems on the switch.

Port Mirroring

The Port Mirroring feature on the switch is very useful for troubleshooting any connection-oriented problem. Any traffic in or out of one or more ports can be mirrored to a single monitoring port to which a network monitor can be attached.

Port Mirroring can be used as a troubleshooting tool or to enhance the security of your network. For example, an Intrusion Detection Service (IDS) server can be connected to the monitor port to detect intruders attacking the network.

As shown in the following figure, port 20 is monitoring ingress traffic (traffic entering the switch) on port 23 and egress traffic (traffic leaving the switch) on port 1. You can attach a device to port 20 to monitor the traffic on ports 23 and 1.

Figure 11 Port Mirroring

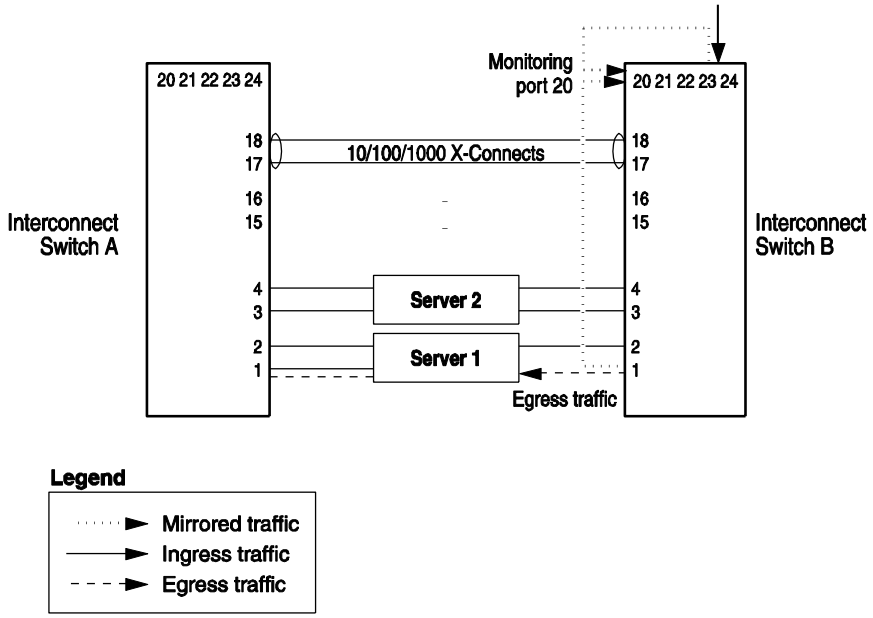

This figure shows two mirrored ports monitored by a single port. Similarly, you can have one mirrored port to one monitored port, or many mirrored ports to one monitored port. The switch does not support a single port being monitored by multiple ports because it supports only one monitored port configured at a time.

Ingress traffic is duplicated and sent to the mirrored port before processing, and egress traffic is duplicated and sent to the mirrored port after processing.

Configuring Port Mirroring (AOS CLI example)

To configure Port Mirroring for the example shown in the preceding figure:

1. Specify the monitoring port.

```
>> # /cfg/pmirr/monport 20              (Select port 20 for monitoring)
```

2. Select the ports that you want to mirror.

```
>> Port 20 # add 23                         (Select port 23 to mirror)

>> Enter port mirror direction [in, out, or both]: in
(Monitor ingress traffic on port 23)

>> Port 20 # add 11                         (Select port 11 to mirror)

>> Enter port mirror direction [in, out, or both]: out
(Monitor egress traffic on port 1)
```

3. Enable Port Mirroring.

```
>> # /cfg/pmirr/mirr ena                     (Enable port mirroring)
```

4. Apply and save the configuration.

```
>> PortMirroring# apply                 (Apply the configuration)
>> PortMirroring# save                   (Save the configuration)
```

5. View the current configuration.

```
>> PortMirroring# cur              (Display the current settings)

Port mirroring is enabled
Monitoring Ports Mirrored Ports
1     none
2     none
3     none
4     none
5     none
   :
   :
17    none
18    none
20(23, in) (11, out)
21    none
   :
```

Configuring Port Mirroring (BBI example)

1. Configure Port Mirroring.
   a. Click the Configure context button.
   b. Open the Switch folder, and select Port-Based Port Mirroring (click the underlined text, not the folder).

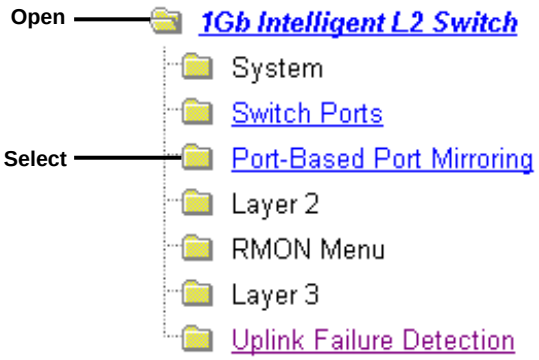

   c. Click a port number to select a monitoring port.

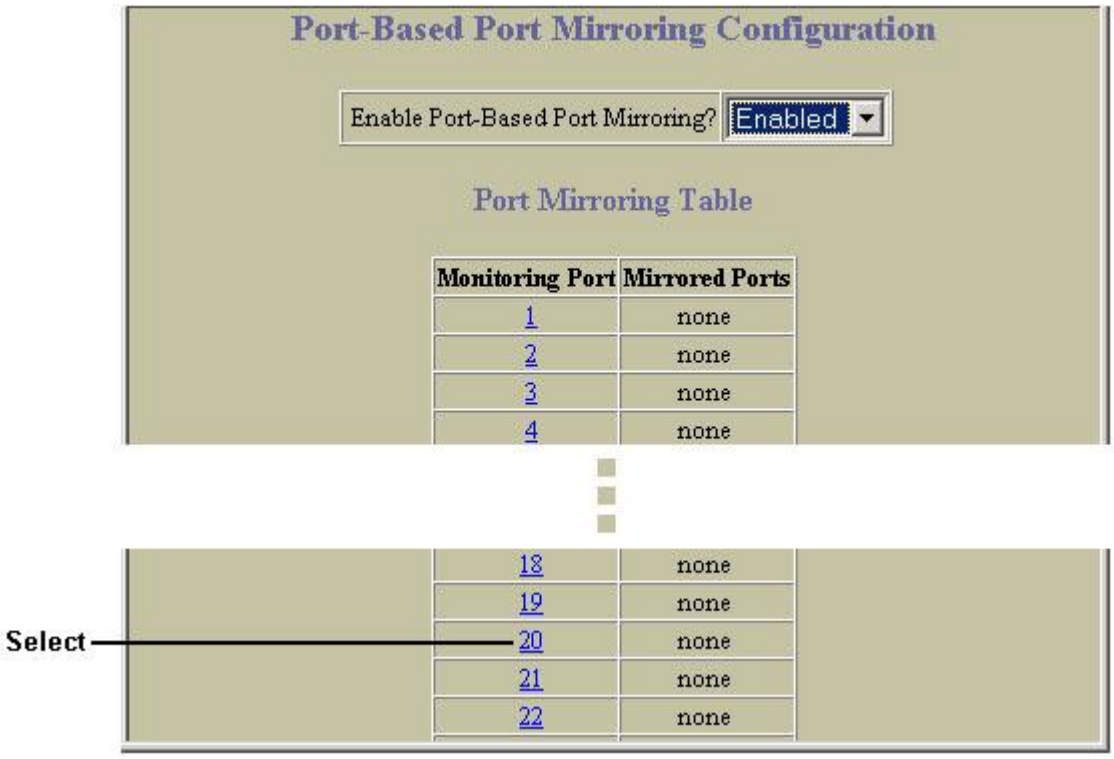

   d. Click Add Mirrored Port.

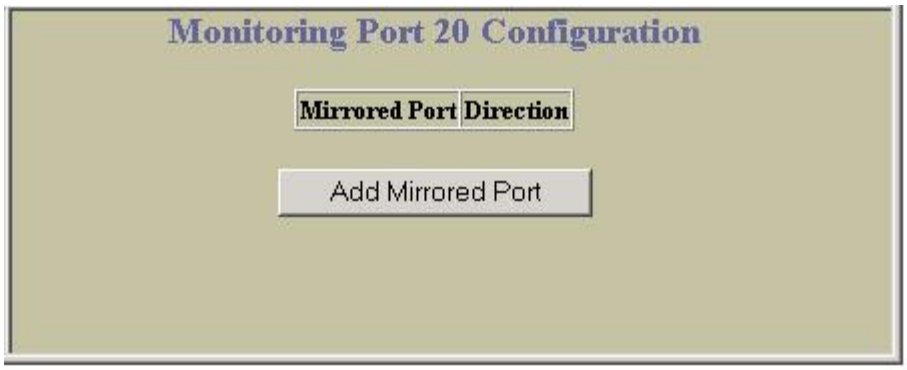

e. Enter a port number for the mirrored port, and select the Port Mirror Direction.

f. Click Submit.

2. Apply, verify, and save the configuration.

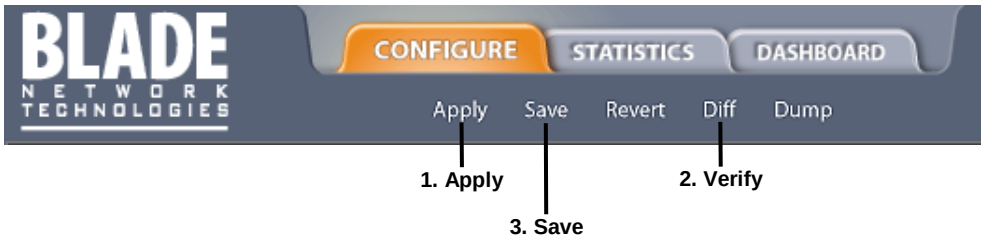

3. Verify the Port Mirroring information on the switch.

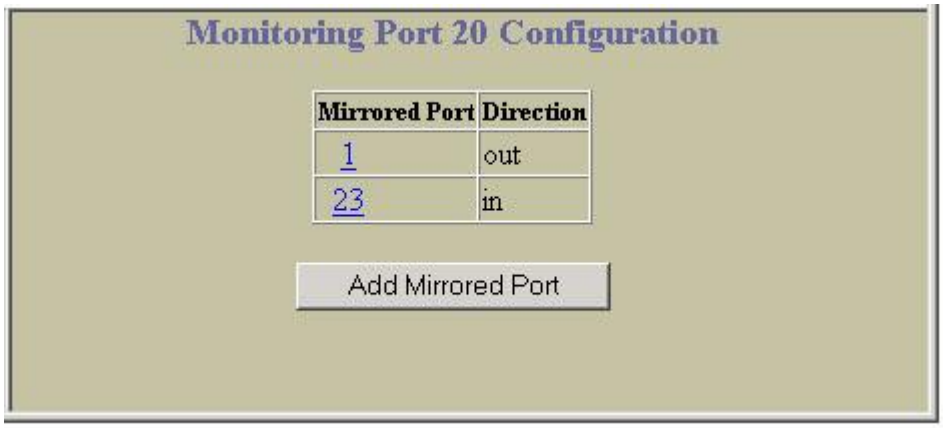

Other network troubleshooting techniques

Other network troubleshooting techniques include the following.

Console and Syslog messages

When a switch experiences a problem, review the console and Syslog messages. The switch displays these informative messages when state changes and system problems occur. Syslog messages can be viewed by using the `/info/sys/log` command. For more information on interpreting syslog messages, see the *Command Reference Guide*.

Ping

To verify station-to-station connectivity across the network, execute the following command:

```
ping <host name> | <IP address> [ (number of tries) [ msec delay ]]
```

The IP address is the hostname or IP address of the device. The number of tries (optional) is the number of attempts (1-32). Msec delay (optional) is the number of milliseconds between attempts.

Trace route

To identify the route used for station-to-station connectivity across the network, execute the following command:

```
traceroute <host name> | <IP address> [<max-hops> [ msec delay ]]
```

The IP address is the hostname or IP address of the target station. Max-hops (optional) is the maximum distance to trace (1-16 devices). Msec delay (optional) is the number of milliseconds to wait for the response.

Statistics and state information

The switch keeps track of a large number of statistics and many of these are error condition counters. The statistics and state information can be very useful when troubleshooting a LAN or Real Server problem. For more information about available statistics, see one of the following:

- "Viewing statistics" chapter of the *Browser-based Interface Reference Guide*, or
- "Statistics Menu" chapter of the *Command Reference Guide (AOS), or*
- “Statistic Commands” chapter of the *Command Reference Guide (ISCLI)*

Customer support tools

The following diagnostics tools are not user-configurable.

- Offline Diagnostics — This tool is used for troubleshooting suspected switch hardware issues. These tests verify that the selected hardware is performing within expected engineering specifications.
- Software Panics — If a fatal software condition is found during runtime, the switch will capture the current hardware and software state information into a panic dump. This dump file can be analyzed post-mortem to determine the cause of the problem.
- Stack Trace — If a fatal software condition occurs, the switch dumps stack trace data to the console.